PCT WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 6 : H04Q 7/38 A2

(11) International Publication Number: WO 00/02406

(43) International Publication Date: 13 January 2000 (13.01.00)

(21) International Application Number: PCT/FI99/00565

(22) International Filing Date: 24 June 1999 (24.06.99)

(30) Priority Data:
981565 7 July 1998 (07.07.98) FI

(71) Applicant (*for all designated States except US*): NOKIA NETWORKS OY [FI/FI]; Keilalahdentie 4, FIN-02150 Espoo (FI).

(72) Inventor; and
(75) Inventor/Applicant (*for US only*): EKBERG, Jan-Erik [FI/FI]; Seljatie 1 A 5, FIN-00320 Helsinki (FI).

(74) Agent: PATENT AGENCY COMPATENT LTD.; Pitkänsillaranta 3 B, FIN-00530 Helsinki (FI).

(81) Designated States: AE, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, US, UZ, VN, YU, ZA, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SL, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

Published
In English translation (filed in Finnish).
Without international search report and to be republished upon receipt of that report.

(54) Title: SYSTEM AND METHOD FOR AUTHENTICATION IN A MOBILE COMMUNICATIONS SYSTEM

(57) Abstract

The invention concerns authentication to be performed in a telecommunications network, especially in an IP network. To allow a simple and smooth authentication of users of IP networks in a geographically large area, the IP network's terminal (TE1) uses a subscriber identity module (SIM) as used in a separate mobile communications system (MN), whereby a response may be determined from the challenge given to the identity module as input. The IP network also includes a special security server (SS), to which a message about a new user is transmitted when a subscriber attaches to the IP network. The subscriber's authentication information containing at least a challenge and a response is fetched from the said mobile communications system to the IP network and authentication is carried out based on the authentication information obtained from the mobile communications system by transmitting the said challenge through the IP network to the terminal, by generating a response from the challenge in the terminal's identity module and by comparing the response with the response received from the mobile communications system. Such a database (DB) may also be used in the system, wherein subscriber-specific authentication information is stored in advance, whereby the information in question need not be fetched from the mobile communications system when a subscriber attaches to the network.

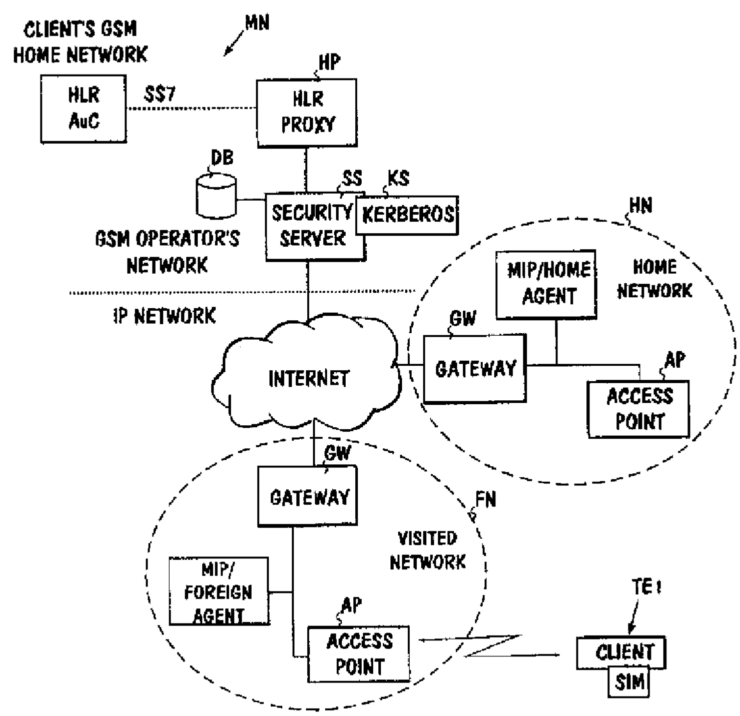

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania
AM	Armenia
AT	Austria
AU	Australia
AZ	Azerbaijan
BA	Bosnia and Herzegovina
BB	Barbados
BE	Belgium
BF	Burkina Faso
BG	Bulgaria
BJ	Benin
BR	Brazil
BY	Belarus
CA	Canada
CF	Central African Republic
CG	Congo
CH	Switzerland
CI	Côte d'Ivoire
CM	Cameroon
CN	China
CU	Cuba
CZ	Czech Republic
DE	Germany
DK	Denmark
EE	Estonia
ES	Spain
FI	Finland
FR	France
GA	Gabon
GB	United Kingdom
GE	Georgia
GH	Ghana
GN	Guinea
GR	Greece
HU	Hungary
IE	Ireland
IL	Israel
IS	Iceland
IT	Italy
JP	Japan
KE	Kenya
KG	Kyrgyzstan
KP	Democratic People's Republic of Korea
KR	Republic of Korea
KZ	Kazakstan
LC	Saint Lucia
LI	Liechtenstein
LK	Sri Lanka
LR	Liberia
LS	Lesotho
LT	Lithuania
LU	Luxembourg
LV	Latvia
MC	Monaco
MD	Republic of Moldova
MG	Madagascar
MK	The former Yugoslav Republic of Macedonia
ML	Mali
MN	Mongolia
MR	Mauritania
MW	Malawi
MX	Mexico
NE	Niger
NL	Netherlands
NO	Norway
NZ	New Zealand
PL	Poland
PT	Portugal
RO	Romania
RU	Russian Federation
SD	Sudan
SE	Sweden
SG	Singapore
SI	Slovenia
SK	Slovakia
SN	Senegal
SZ	Swaziland
TD	Chad
TG	Togo
TJ	Tajikistan
TM	Turkmenistan
TR	Turkey
TT	Trinidad and Tobago
UA	Ukraine
UG	Uganda
US	United States of America
UZ	Uzbekistan
VN	Viet Nam
YU	Yugoslavia
ZW	Zimbabwe

SYSTEM AND METHOD FOR AUTHENTICATION IN A MOBILE COMMUNICATIONS SYSTEM

Field of the invention

The invention relates to authentication in a telecommunications network, especially in an IP network (IP = Internet Protocol), and also to improvement of the network's data security features with the aid of the performed authentication. Authentication means verification of the identity of the party, such as the subscriber, who has generated data. Using authentication it is also possible to guarantee integrity and confidentiality of the said data. Authentication may be performed for various purposes, such as for checking the right of use of network services. The invention is intended for use especially in connection with mobile terminals, but with the solution according to the invention advantages are also achieved in connection with fixed terminals.

Background of the invention

The strong growth in number of Internet users has been one of the most remarkable phenomena in communications in recent years. The rapid growth has also highlighted defects on the Internet. One of these is the poor data security of the network. The IP protocol version (IPv4) now in general use does not provide any such means, with which it would be possible to make sure that information arrived from the opposite end did not change during the transfer or that the information did in fact arrive from that source, who claims to have sent the information in question. In addition, it is easy to use various tools in the network for listening in to the traffic. For these reasons, those systems are very vulnerable which transmit non-encrypted critical information, e.g. passwords.

The new IP version (IPv6) has internal characteristics that allow safe communication between Internet users. Because the transition to the new protocol will be slow, the data security features should be such that they are compatible with the present IP version (IPv4), and so that they can be added to this.

Various such systems have been developed to improve the data security properties of the Internet where users can send the information encrypted to the other party. One such system is the Kerberos, which is a service with which network users and services can authenticate one another and with which users and services can bring about encrypted connections between each other. The Kerberos system is utilised in one embodiment of the present

invention which will be described more closely hereinafter.

Another current trend is the strongly increasing use of various mobile terminals. Along with this trend it is even more important that the terminals will have access to the data network also when being located outside their own home network. Such an access can essentially improve the usability of e.g. a portable computer, when the user is not in his/her usual working environment. Points of access may be located e.g. at airports, in railway stations, in shopping malls or on any other public premises, and the access may be wired or wireless.

Systems of the described kind, which can be used for sending encrypted information between parties, are mainly intended for fixed terminals and they require that the users are registered in advance as users of the service. It is a problem nowadays that for IP networks supporting mobility of the terminals there is no such existing and functioning authentication or key management system that would guarantee good geographical coverage and at the same time allow the user easily to have an authenticated and safe connection available to himself/herself in an area which is geographically as large as possible.

Summary of the invention

It is a purpose of the invention to eliminate the drawback described above and to bring about a solution, with which users of a telecommunications network, such as an IP network, can be simply and smoothly authenticated, almost irrespectively of where their network access point is located geographically at each time.

This objective is achieved through the solution defined in the independent claims.

The invention utilizes the authentication method of an existing mobile communications network, especially the GSM network (Global System for Mobile Communications), in an IP network (or in any other network which is separate from the mobile communications network). This means that a user of the IP network in his IP network terminal uses the same (or an essentially similar) subscriber identification unit (SIM) as in his mobile phone or station. The idea is to fetch the subscriber's authentication data from the mobile communications network over to the IP network side and to carry out the authentication in the IP network based on this data. The mobile network is not necessarily a GSM network, but it may be some other mobile communications net-

work, wherein authentication is used essentially in the same manner, e.g. a DCS network (Digital Cellular System), a GPRS network (General Packet Radio Service, which is a sub-network of the GSM) or a UMTS network (Universal Mobile Telecommunications System).

In an advantageous embodiment of the invention, the user is registered in response to a successful authentication into a separate key management system, preferably a Kerberos system, whereby it is possible then easily to bring about an encrypted channel between users communicating with one another. This is especially important when at least a part of the transmission path consists of a radio path.

Owing to the solution according to the invention, users of the IP network are easily and smoothly authenticated and, in addition, the users are able to avail themselves of efficient security features in a geographically large area. This is due both to the widespread use of GSM networks and to the fact that roaming agreements between operators allow authentication of subscribers entering a foreign network. E.g. today (1998) a Finnish GSM operator has common traffic agreements with operators working in more than 60 countries.

Owing to the solution according to the invention, ISP (Internet Service Provider) operators typically also providing mobile communication services need not separately procure authentication and key management systems in the IP network, but they may use also for this purpose the features of the mobile communications network which they operate.

With the solution according to the invention such an advantage is also achieved in connection with fixed terminals, that functions built in connection with the mobile communications network can be utilised in connection with Internet services. E.g. an organisation working both as a mobile communication operator and as an ISP operator may use charging services built in connection with the mobile communications network for charging for the Internet services which he provides. When also fixed terminals are authenticated with the method according to the invention, much certainty is achieved that the bill will be directed at the correct subscriber. In addition, the subscriber can be authenticated, even if he attaches to the network from a foreign terminal.

A brief description of the drawings

In the following, the invention and its preferred embodiments will be described more closely referring to the examples shown in Figures 1...10 in the appended drawings, wherein

Figure 1 illustrates an operating environment of the method in accordance with the invention,

Figure 2 shows an exchange of messages between various elements, when the terminal attaches to the network or detaches from the network,

Figure 3 illustrates the structure of those messages, with which the server of the system is told that the user has attached to the network or has detached from the network,

Figure 4 shows an exchange of messages taking place between the various elements during authentication,

Figure 5 illustrates the general structure of the messages shown in Figure 5,

Figure 6 illustrates those elements of the system, which are used for acquiring a connection-specific encryption key between two terminals,

Figure 7 shows an exchange of messages taking place in order to obtain an initial ticket from the Kerberos server,

Figure 8 illustrates those parts of a terminal which are essential from the view-point of the invention,

Figure 9 shows an exchange of messages taking place in order to obtain an encryption key for communication between two terminals, and

Figure 10 illustrates an alternative embodiment of the system.

Detailed description of the invention

In the following the invention will be described with reference to a network environment, wherein mobility of the subscribers is supported with the aid of a Mobile IP protocol (MIP hereinafter). The MIP is such a version of the existing IP, which supports mobility of the terminals. (The MIP principle is described e.g. in the RFC 2002, October 1996, or in the article Upkar Varsh-ney, *Supporting Mobility with Wireless ATM*, Internet Watch, January 1997.)

The MIP is based on the idea that each mobile host or mobile node has an agent (home agent) allocated for itself, which relays packets to the current location of the mobile node. When the mobile node moves from one sub-network into another, it registers with the agent (foreign agent) serving the concerned sub-network. The last-mentioned performs checks with the mobile node's home agent, registers the mobile node and sends the registration information to it. Packets addressed to the mobile node are sent to the mobile node's original location (to the home agent), thence they are relayed further to the current foreign agent, which will forward them to the mobile node.

Figure 1 shows a typical operating environment of the method in accordance with the invention. The heart of the system is the security server SS, which is connected both to the Internet and to a proxy server HP, which has access to a separate mobile network MN, which in this example is a GSM network. The proxy server forms a network element, which (in a manner to be described later) relays traffic between the security server and the home location registers HLR of mobile communications networks, which home location registers HLR are located in the home networks of the subscribers. In practice, both the proxy server and the security server are located on the premises of the network operator, e.g. in the same room, so that even if there is an IP connection between the security server and the proxy server, it is a secured connection. As the GSM network is known as such and the invention does not require any changes to be made in it, it is not described more closely in this connection.

Users moving in the area of the system can use portable computers, PDA equipment, intelligent phones or other such terminals. Only one terminal TE1 is illustrated by reference mark CLIENT in the figure. For the present purposes, client generally means an object using the services provided by the network and carried out by the network servers. Client often means a program which connects with a server on behalf of the network user.

Two sub-networks are shown in the figure and in practice they may be e.g. Ethernet local area networks, wherein TCP/IP packets are transmitted: the user's home network HN and the foreign network FN, to which terminal TE1 is assumed to be connected. These sub-networks are both connected to the Internet by way of a gateway GW (a router). The home network includes the home agent HA of the said mobile host and the foreign network correspondingly includes the foreign agent FA. Accesses to the sub-networks take place through access points AP, e.g. in a wireless manner, as is shown in the figure.

The terminals are formed by two parts in the same way as the ordinary GSM telephone: of the subscriber device proper, e.g. a portable computer (with software) and of the SIM (Subscriber Identity Module), whereby from the viewpoint of the network the subscriber device becomes a functioning terminal only when the SIM has been pushed into it. In this case described as an example, the SIM is the subscriber identity module for use in the GSM network. A terminal may have access only to the IP network, or it may be a so-called dual mode device, which has access both to the IP network and to the GSM network. The access to the IP network takes place e.g. with the aid of a LAN card

in the terminal and to the GSM network with the aid of a GSM card, which in practice is a stripped telephone, which is located e.g. in the computer's PCMCIA expansion slot.

In a preferred embodiment of the invention, there is also a Kerberos server KS in connection with the security server which is known as such and which is used for implementing encrypted connections in a manner to be described hereinafter. The security server and the Kerberos server may be physically in the same machine.

For the security server to know when the user enters or exits the IP network, a channel is brought about between the security server and the home agent in the manner shown in Figure 2. In accordance with the MIP protocol, foreign agent FA continuously sends broadcast messages to its own sub-network, which messages are called by the name of "agent advertisement" and which are indicated by the reference mark AA in the figure. When the terminal attaches to the said sub-network, it will receive these messages and conclude from them whether it is in its own home network or in some other network. If the terminal finds that it is in its home network, it will function without any mobility services. Otherwise the terminal will get a care-of address in the foreign network in question. This address is the address of that point in the network to which the terminal is temporarily connected. This address at the same time forms the termination point of the tunnel leading to the said terminal. Typically, the terminal gets the address e.g. from the above-mentioned broadcast messages, which the foreign agent is sending. Thereupon the terminal sends a RR (Registration Request) to its own home agent through foreign agent FA. The message contains, among other things, that care-of address, which the terminal just received. Based on its received request message, the home agent updates the said terminal's location information in its database and through the foreign agent it sends a Registration Reply R_Reply to the terminal. In the reply message there is all the necessary information indicating how (on what conditions) the home agent has accepted the registration request.

All the messages between the terminal, the foreign agent and the home agent which were described above are normal messages in accordance with the MIP protocol. The mobile node may also register directly with the home agent. The above-mentioned RFC describes the rules, which determine whether the mobile node will register directly with the home agent or through the foreign agent. If the mobile node gets a care-of address in the manner

described above, the registration must always be made through the foreign agent. According to the MIP protocol, authentication is also performed in connection with the registration with the purpose to reduce the occurrence of errors in connection with the registration. The registration is based on a check value calculated from the registration message (from the registration request or reply), and the registration must be made only between that mobile node and that home agent, which have a shared fixed key (which is agreed upon in advance). Under these circumstances, the foreign agent is not necessarily able to authenticate the mobile node. This problem is aggravated, if as large a geographical coverage as possible is an objective in the system.

According to the invention, a facility is added to the home agent to the effect that the home agent provides the security server with information about the terminal attached to the network, after the registration request message has arrived from the foreign agent. This message is indicated in the figure by reference mark MOB_ATTACH. Correspondingly, the home agent provides the security server with information about the terminal which has left the network after the terminal has detached from the network (after the terminal has detached from the network or after the lifetime of the address given to it has run out). In the figure, this message is indicated by the reference mark MOB_DETACH. To each type of message the security server sends an acknowledgement message (MOB_ACK). As regards their purpose of use, the MOB_ATTACH and MOB_DETACH messages correspond to the IMSI attach/detach procedures used in a GSM network.

The home agent monitors the replies arriving from the security server and sends the messages again (with the same parameters), should no acknowledgement message arrive from the security server within a predetermined time, e.g. 30 seconds.

Figure 3 illustrates the structure of the MOB_ATTACH, MOB_DETACH and MOB_ACK messages. In the messages there is a type field 31, which identifies the type of the message, a number field 32, which contains the random number or sequence number identifying the session, and an address field 33, which contains the client's IP address. The last-mentioned field is absent from the acknowledgement message. The messages are transmitted in fields reserved for the payloads of IP datagrams.

Thus, when the terminal has attached to the network, the security server receives from the home agent information about the IP address of the concerned terminal. Thereupon follows authentication of the client, which will

be described in the following with reference to Figure 4. For the authentication, the security server first asks the client for the IMSI (International Mobile Subscriber Identity), which is stored on the SIM (the AUTH_ID_REQ message). To this the client replies by giving his IMSI (which is a 9-byte identifier in accordance with the GSM specification) in the AUTH_ID_RSP reply message. The inquiry travels through the home agent to the termination point of the above-mentioned tunnel, but the reply comes directly from the terminal to the security server.

If the client's IP address does not change often, it is preferable to store in the security server the IMSI identifiers corresponding to the IP addresses, whereby identifiers need not be moved around unnecessarily in the network. Thus, the above-mentioned messages are not necessary.

When the terminal has stated its IMSI identifier or when the security server has fetched it from its database, the security server starts the actual authentication. To enable authentication of the terminal's SIM, there must be a connection between the security server and the AuC (Authentication Center) located in connection with the home location register HLR of the subscriber's own GSM network. This is implemented with a proxy server HP, which functions as a connecting network element between the IP network and the GSM network, more precisely between the IP network and the SS7 signaling network utilized by the GSM network. The GSM network service needed in the authentication is MAP_SEND_AUTHENTICATION_INFO (GSM 9.02, v. 4.8.0). This service is implemented by using the proxy server HP, which may be located on the premises of the local GSM operator. The security server transmits to the proxy server a SEC_INFO_REQ authentication request message, which contains a session identifier and the IMSI subscriber identifier. The proxy server for its part transmits to the authentication centre AuC an inquiry message in accordance with the MAP (Mobile Application Part) protocol, which inquiry message is used to request an authentication triplet and which is normally transmitted between the VLR and the HLR. In response to this inquiry message, the HLR returns to the proxy server a normal authentication triplet, which contains a challenge (RAND), a response SRES (Signed Response) and a key Kc (the connection-specific encryption key used in the GSM network). The proxy server relays the triplet further to the security server in a SEC_INFO_RSP message. The security server stores the triplet and transmits the challenge (the AUTH_CHALLENGE_REQ message) further to the terminal's SIM, which based on this message generates a response

(SRES) and a key Kc. The terminal stores the key and transmits the response (the AUTH_CHALLENGE_RSP message) (SRES) back to the security server.

In the terminal there is preferably a database, wherein the challenges are stored. In this way it is possible to make sure that one challenge will be used just once. In this manner it is possible to prevent anyone from pretending to be a security server by snatching from the network the (non-encrypted) challenge and the response and by finding out the key Kc from these. If the same challenge occurs once again, no reply will be given to this challenge. The security server may also filter out those challenges which have already been used, and when required it may ask for a new authentication triplet from the GSM network, so that no such challenge which has already been used will be transmitted to the terminal.

The proxy server HP functions in the system as a virtual visitor location register VLR, because at least as regards the authentication triplet inquiries it appears from the home register like a network element of the same kind as the genuine visitor registers of the GSM network. The proxy server also functions as a filter allowing access to the GSM system's signaling network only to authentication triplet inquiries. The proxy server does not either interfere with any other inquiries from the home register on the GSM network side.

Figure 5 illustrates the general structure of the messages presented in Figure 4. In the messages there is a type field 51, which identifies the type of the message, a number field 52, which contains the random number or sequence number identifying the session, and a payload field 53, the length of which varies depending on which message is at issue. In messages between the security server and the terminal, the two first fields occur in all messages, but there is no payload field in the AUTH_ID_REQ message. In the AUTH_ID_RSP message the length of the payload field is 9 bytes (the length of IMSI is 1+8 bytes), in the AUTH_CHALLENGE_REQ message its length is 16 bytes (the length of RAND is 16 bytes) and in the AUTH_CHALLENGE_RSP message its length is 4 bytes (the length of SRES is 4 bytes). In the messages between the security server and the proxy server, the length of the payload field is 9 bytes (IMSI) in the case of the SEC_INFO_REQ message and n×28 bytes in the case of the SEC_INFO_RSP message (in the triplet there is a total of 28 bytes and the network elements are generally configured so that they will transmit 1...3 subscriber-specific triplets at a time). As mentioned above, normal GSM network signaling is used between the proxy server and the home location regis-

ter HLR.

The security server compares the response it received from the terminal with the response arrived in the triplet and, if it is found in the comparison that the responses are the same, the authentication is successful.

In response to a successful authentication, the security server starts a registration with the Kerberos server. In this context the Kerberos server means a process, which provides a Kerberos service. The Kerberos server is preferably located in connection with the security server, as is shown in Figure 1.

Kerberos is a system intended for authentication of network users and services. It is a trusted service in the sense that its every client trusts that the system's assessment of all its other clients is correct. Since the Kerberos system is known as such, and its operation is not changed in any way, it will not be described in detail in this context. The system is described e.g. in the document Steiner, Neuman, Schiller: Kerberos: An Authentication Service for Open Network Systems, January 12, 1988, from which the interested reader may find background information, if he so desires. In the following description the same ways of marking will be used as in the above-mentioned document. The description is based on the Kerberos version 4.

c	$\rightarrow$	client,
s	$\rightarrow$	server
c-addr	$\rightarrow$	client's network address
tgs	$\rightarrow$	ticket-granting server
K_x	$\rightarrow$	x's private key
$K_{x,y}$	$\rightarrow$	session key for x and y
$\{abc\}K_x$	$\rightarrow$	abc encrypted using x's personal key
Tx,y	$\rightarrow$	x's ticket for using y.

Figure 6 illustrates the objects of the Kerberos and authentication applications. It is assumed in the figure that the system has two clients, A and B. Each client may be a terminal, which has been authenticated by the security server in the manner described above, when it attached to the IP network, or one may be a “permanently” authenticated client, e.g. a server. The Kerberos application includes two parts: client program KC, which is located at the terminal, and server program KS, which is located at the security server. The server program also includes a ticket-granting server TGS. Correspondingly, the authentication application includes two parts: the client program AC, which

is located at the terminal, and the server program AS, which is located at the security server. Communication takes place with the aid of IP/MIP/IP-SEC stacks, which will be described in greater detail below.

The following is a description of how the Kerberos protocol is used for bringing about a connection-specific key between terminals A and B.

When the security server has found that the authentication was successful, it will start registration of the Kerberos client with the Kerberos server. In practice, this happens in such a way that the security server's authentication block AS registers the key Kc arrived in the authentication triplet (a) as the client's password and (b) as a password into the service formed for the client's IP address or for the IMSI subscriber identifier. The service is given some name which is determined in advance.

Then the client may request a ticket for the ticket-granting server using the key Kc. This exchange of messages is shown in Figure 7. After the client has received the key Kc, it transmits to the security server (to the Kerberos server) a message, with which it requests an initial ticket of the Kerberos system. There may be a brief predetermined delay between the reception of the key and the transmission of the message, so that the security server will have time first to perform the registration with the Kerberos server. After the delay, the terminal transmits to the security server a request in accordance with the Kerberos protocol, which always contains the client's identity (the IMSI or IP address) and the name tgs of a certain special service, the ticket-granting service. Upon receiving this inquiry the Kerberos server checks whether it knows the client. If it does, it will generate a random connection-specific key $K_{c,tgs}$, which will be used later in data transmission between the client and the ticket-granting server. Thereupon the Kerberos server generates a ticket $T_{c,tgs}$, with which the client may use the ticket-granting service. This ticket contains the client's name, the name of the ticket-granting server, the current time of day, the lifetime of the ticket, the client's IP address and the connection-specific key just generated. Using the methods of marking described above, the contents of the ticket can be presented as follows $T_{c,tgs}=\{c, tgs, timestamp, lifetime, c\text{-}addr, K_{c,tgs}\}$. This ticket is encrypted using key K_{tgs}, which is known only to the ticket-granting server and to the Kerberos server. Then the Kerberos server transmits as a response to the client a packet, which contains the encrypted ticket and a copy of the connection-specific key $K_{c,tgs}$. The response is encrypted using the client's own key Kc. The terminal stores

the ticket and the session key for future use.

When the terminal has stored the ticket and the session key, it has access during the ticket's lifetime to the ticket-granting service and it is prepared to be in connection with a third party.

Figure 8 illustrates those functional blocks of a terminal, which are essential from the viewpoint of the invention. The terminal is in connection with the network by way of the IP/MIP/IP-SEC protocol stack. IP/MIP/IP-SEC is such a known TCP/IP stack, which has built-in mobile IP characteristics and encryption functions. Seen from above, this stack appears just like an ordinary IP stack, but from below (from the network side) the said stack transmits encrypted information in accordance with a certain security policy. This security policy is determined by a separate security policy block SPB, which controls the IP/MIP/IP-SEC stack by indicating to the stack the other objects in the network to which encrypted information must be sent. These objects are generally defined in the security policy block with the aid of the terminal's IP address and port number. The definition can be made even finer by also defining those user identifiers, for which the encryption is done. In practice, the security policy block is built into the IP/MIP/IP-SEC stack, but in a functional sense it is a block in its own right.

In addition to the security policy block, the terminal contains a key management block KM, which attends to management of keys. In connection with the key management block there is a database containing all the encryption keys used by the terminal. The key management block can be implemented e.g. with the aid of the known PF_KEY API (API=Application Programming Interface). PF_KEY is a generic application programming interface, which may be used not only for IP layer security services, but also for other security services of the network. This API determines the socket protocol family, which the key management applications use to communicate with parts of the operating system relating to the key management. Since the invention is not related to the known PF_KEY protocol, it will not be described more closely in this context. The protocol is described in the document McDonald, Metz, Phan: PF_KEY Management API, version 2, 21 April, 1997, where the interested reader will find background information.

In the key management block KM there are specific definitions for how and with which key the encryption is carried out to each network address. This definition may be made e.g. so that for each individual IP address and port that protocol and that key are stated which must be used when in connec-

tion with the port in question.

When a packet which is to be transmitted outwards arrives in the IP/MIP/IP-SEC stack, the stack reads the packet's destination address and asks the security policy block SPB which is the encryption policy as regards a packet carrying the address in question. In response, the security policy block tells the IP/MIP/IP-SEC stack whether encryption is to be made, and if so, with which method the encryption is to be carried out. This information is relayed to the key management block KM.

In the initial stage, the user has determined those connections for the security policy block, on which encryption must be used. If the security policy block states that encryption must be used and if the key management block finds that there is as yet no key for the terminal with which a connection is desired, the key management block will send a key request to the Kerberos client KC, who will request a server ticket for the concerned terminal from the security server's ticket-granting service. This signalling is illustrated in Figure 9. The terminal (the Kerberos client) sends to the ticket-granting server such a request in accordance with the Kerberos protocol, which contains the name (s, e.g. terminal B) of that server, for which the ticket is desired, a ticket $T_{c,tgs}$ encrypted with the ticket granting server's own key K_{tgs} for access to the ticket-granting service and an authenticator Ac, which is encrypted with a connection-specific key $K_{c,tgs}$. The authenticator is a data structure, which contains the client's name and IP address as well as the current time. Observing the used method of marking Ac = {c, c-addr, timestamp}.

The ticket-granting server checks the authenticator's information and the ticket $T_{c,tgs}$. If the ticket is all right, the ticket-granting server generates a new random session key $K_{c,s}$, which the client may use together with a third party of his choice. Then the ticket-granting server forms a new ticket $T_{c,s}$ for the said third party, encrypts the ticket using the said third party's own key K_s, which is the same as the concerned subscriber's key Kc described above, and transmits the encrypted key together with the session key to the terminal. The entire reply is encrypted using key $K_{c,tgs}$.

Upon receiving the reply message, the terminal unpacks the packet, transmits the first part $\{T_{c,s}\}K_s$ to the third party (to terminal B) and stores the new session key $K_{c,s}$ in the key database. The terminal of the third party gets the recently generated session key $K_{c,s}$ from the ticket by first decrypting the ticket with its own key Kc. Thereafter the new session key is available to both

terminals and encrypted data transmission may begin.

When the Kerberos client has started his activity (when the client is registered with the Kerberos server), it must inform the IP/MIP/IP-SEC layer that it is able to serve session key requests. By using the PF_KEY protocol, this is done in such a way that the Kerberos client opens a special socket address into the kernel of the operating system and registers with the kernel with a SADB_REGISTER message. Then the PF_KEY protocol sends a SADB_ACQUIRE message each time when the key is needed for some outbound interface. When receiving this message, the Kerberos client will act in the manner described above, that is, he sends a request to the ticket-granting server, of the received response it sends the part intended for the other party to the opposite end of the connection and relays the received session key to the key management block. In addition, the Kerberos client listens to a certain socket address in order to notice any tickets that may arrive from other objects in the network. Having received such a ticket packet, it acknowledges reception of the packet, unpacks the packet and relays the necessary keys to the key management system, whereby these keys can be used when connections exist with the concerned peer.

When the terminal detaches from the network (message MOB_DETACH), the security server will remove both registrations from the Kerberos server.

In practice, the terminal and the security server must have certain port numbers open for non-encrypted data transmission. Such ports are the port, through which authentication messages are transmitted between the terminal and the server (Figure 4), the port, through which tickets are transferred to the Kerberos clients, and the port, through which ticket requests are transferred.

The authentication triplet can be sought in various ways. In a small-scale embodiment it is possible to use a virtual "HLR database", wherein a suitable number of authentication triplets is stored in advance. E.g. 10000 triplets from each user would require 280 kilobytes of memory per user. Thus, e.g. a 6 GB disk could accommodate authentication triplets for more than 21000 users. The authentication triplets may be loaded in advance when the user gets the service, by leaving the SIM module for a few hours in a smart card reader, which supplies the challenges to the module. The authentication triplets formed of the obtained responses are stored in the database using the module's information. This method also works with all SIM modules, irrespective of the operators. The database may be located e.g. in connection with the

security server. Thus, it is not necessary to seek the authentication triplet(s) from the mobile communications network, but subscriber-specific authentication triplets can be stored in advance in a database DB located in connection with the security server (compare with Figure 1). This means that proxy servers are not necessarily needed at all. For some subscribers there may also be ready-made authentication triplets in the database and for some they may be fetched in real time from the mobile communications system. Authentication triplets can also be fetched in advance from the mobile communications system and placed in the database.

In principle, it is also possible to copy each user's SIM module and use the copy in connection with the security server for authentication of the user (whereby no inquiry is made from the mobile communications network).

These two methods described above make it possible for the used SIM modules to be modules dedicated solely for this purpose, and they do not necessarily relate to the mobile communications network's subscriber.

The necessary authentication data can also be obtained from the GSM network e.g. from the connection between the MSC (Mobile Switching Centre) and the BSC (Base Station Controller). Thus, the proxy server need not necessarily emulate the visitor location register VLR, as was presented above, but it may also function as a network element of the same kind as the GSM network's base station controller. Such an alternative is illustrated in Figure 10, where the said network element is marked with the reference mark BP. In this case, the proxy server is thus a virtual base station controller, which is connected to the MSC (Mobile Switching Centre) in the same way as the GSM network's normal BSCs (Base Station Controllers). Looking from the mobile switching centre, the proxy server looks like an ordinary base station controller at least as regards the signalling relating to authentication.

However, it is a problem in this second alternative that it requires considerably more complex signalling between the proxy server and the GSM network than the first alternative (Figure 1). Besides, in consequence of the authentication of the second alternative, the user will in the GSM system move into the area of the proxy server BP emulating a base station controller, but this is not a real base station controller in the sense that it would be able also to switch calls. Thus, this solution can be used only in connection with data services, and the terminal can not be the kind of dual mode equipment as mentioned above.

Although the invention was described in the foregoing with reference

to a MIP enabled network, the solution according to the invention is not bound to this protocol. If the protocol to be used is IPv6, then there are no proper agents in the network. Hereby the information about when the user is in the network must be sought from the routing tables of the router in the user's home network. In practice, this means that the network must include a separate "locating agent", which by monitoring or "pinging" the router will notice that the user has entered the network and in consequence of this will start authentication by sending to the security server a message (MOB_ATTACH) about the new user. It is probable, however, that router manufacturers are designing a protocol from which it emerges when the user is in the network.

Although the invention was described above with reference to the examples shown in the appended drawings, it is obvious that the invention is not limited to these, but it may be modified within the inventive idea presented in the appended claims. Authentication need not necessarily be performed in order to set up an encrypted connection between users, but as a result of a successful authentication one may perform e.g. registration with a mail server before transmitting e-mail messages to the user's machine. In this way a more reliable authentication is achieved than by the present methods based on passwords. In addition, in connection with the access points there may be local servers, which function as proxy servers for the security server proper, or the system may include more than one security server. Instead of the Kerberos system it is also possible to use e.g. public key management, which is based on a x.500-database and on x.509 certificates.

Claims

1. Authentication method for telecommunications networks, especially for IP networks, in accordance with which method the identity of a subscriber attached to the network is authenticated,

characterized by

- in a network terminal (TE1), using a subscriber identity module (SIM) essentially of the same kind as in a known mobile communications system (MN), which identity module is such that a response is obtained as a result of a challenge given to it as input,

- using a special security server (SS) in the network so that when a terminal attaches to the network, a message of a new user is transmitted to the security server,

- fetching subscriber authentication information corresponding to the said new user from the said mobile communications system to the said network, which authentication information contains at least a challenge and a response, and

- performing the authentication based on the authentication information obtained from the mobile communications system by transmitting the said challenge to the terminal through the network, by generating a response from the challenge in the identity module of the terminal and by comparing the response with the response received from the mobile communications system.

2. Method as defined in claim 1, characterized in that fetching of the subscriber's authentication information from the mobile communications system is started from the security server (SS) in response to the said message.

3. Method as defined in claim 1, characterized in that in response to a successful authentication, registration of the subscriber is performed as a client of a separate key management system.

4. Method as defined in claim 3 for IP networks, characterized in that the known Kerberos system is used as the key management system.

5. Method as defined in claim 4, characterized in that the subscriber-specific authentication information obtained from the mobile communications system also includes a key (Kc), whereby the subscriber is registered as a client of the Kerberos system so that the key is registered (a) as the client's password and (b) as a password for a service formed for the client's IP address or for a subscriber identity (IMSI) used in the mobile communications system.

6. Method as defined in claim 1, characterized in that the subscriber's authentication information is fetched with the aid of a separate proxy server (HP), which functions as a network element emulating the visitor location register VLR of the mobile communications system and which requests the authentication information from an authentication centre AuC located in connection with the subscriber's home location register HLR in the same way as the mobile communications system's own visitor location register.

7. Method as defined in claim 1, characterized in that the subscriber's authentication information is fetched with the aid of a separate proxy server (BP), which functions as a network element emulating the mobile communications system's base station controller and which is in connection with the mobile communications system's mobile switching centre (MSC) for fetching the authentication information from an authentication centre AuC located in connection with the subscriber's home location register HLR in the same way as the authentication information is fetched to the mobile communications system's own base station controller.

8. Authentication system for telecommunications networks, especially for IP networks, which system includes authentication means for authenticating the identity of a subscriber who has attached to the network,

characterized in that the authentication means include

- a subscriber identity module (SIM) connected to the network's terminal (TE1), the module being essentially similar to the subscriber identity module used in a separate mobile communications system (MN), whereby a response can be determined from a challenge given to the identity module as input,

- messaging means (HA) for sending a message when a terminal attaches to the network,

- a special security server (SS) for receiving the said message,

- means for requesting authentication information corresponding to a subscriber from the said mobile communications system (MN), which information contains at least a challenge and a response, and

- on the side of the said network, data transmission and checking means for transmitting the challenge through the network to the identity module, for returning the response from the terminal to the network and for comparing the received response with the response received from the mobile communications system.

9. System as defined in claim 8, characterized in that the said identity module is the subscriber identity module (SIM) used in the GSM network

10. System as defined in claim 8, characterized in that the messaging means are adapted into a home agent (HA) in accordance with the mobile IP network.

11. System as defined in claim 8, characterized in that the means for requesting authentication information include the said security server and a proxy server (HP, BP), which is connected to the GSM network.

12. System as defined in claim 11, characterized in that the proxy server functions as a network element emulating the visitor location register VLR of the GSM network.

13. System as defined in claim 11, characterized in that the proxy server functions as a network element emulating the base station controller BSC of the GSM network.

14. System as defined in claim 11, characterized in that the system further includes a Kerberos server (KS) which is known as such and as the user of which the subscriber will be registered as a result of a successful authentication.

15. Authentication method for telecommunications networks, especially for IP networks, in accordance with which method the identity of a subscriber attached to the network is authenticated,

characterized by

- in a network terminal (TE1), using a subscriber identity module (SIM) essentially similar to the one used in a known mobile communications system (MN), which identity module is such that a response is obtained as a result of a challenge given to it as input,

- storing subscriber-specific authentication information in a database (DB), the information being in that way essentially similar to the information used for authentication in the said mobile communications system that it contains at least a challenge and a response,

- using a special security server (SS) in the network so that when a terminal attaches to the network, a message about the new user is transmitted to the security server,

- in response to the message, retrieving authentication information of the subscriber corresponding to the new user from the said database (DB), and

- performing authentication based on the authentication information obtained from the database by transmitting the said challenge through the network to the terminal, by generating a response from the challenge in the identity module of the terminal and by comparing the response with the response obtained from the database.

16. Method as defined in claim 15, characterized in that the database is stored in connection with the security server.

17. Method as defined in claim 15, characterized in that in response to a successful authentication, registration of the subscriber is performed as the user of a separate key management system.

18. Method as defined in claim 17, characterized in that the known Kerberos system is used as the key management system.

19. Authentication system for telecommunications networks, especially for IP networks, which system includes authentication means for authentication of the identity of a subscriber attached to the network,

characterized in that the authentication means include

- a subscriber identity module (SIM), which is connected to a network terminal (TE1) and which is essentially similar to the subscriber identity module used in a separate mobile communications system (MN), whereby a response can be determined from the challenge given as input to the identity module,

- messaging means (HA) for sending a message when a terminal attaches to the network,

- a special security server (SS) for receiving the said message,

- database means (SS, DB), which include a database (DB), wherein subscriber-specific authentication information is stored, which is in such a way essentially similar to the information used for authentication in the said mobile communications system that it includes at least a challenge and a response, and retrieval means (SS) for retrieving subscriber-specific authentication information from the said database in response to the message,

- on the side of the said network, data transmission and checking means for transmitting the said challenge through the network to the identity module, for returning the response from the terminal to the network and for comparing the received response with the response received from the database.

20. System as defined in claim 19, characterized in that the said identity module is a subscriber identity module (SIM) used in the GSM network.

21. System as defined in claim 19, characterized in that the messaging means are adapted into a home agent (HA) in accordance with the mobile IP network.

22. System as defined in claim 19, characterized in that the system further includes a Kerberos server (KS), which is known as such and as the client of which the subscriber is registered as the result of a successful authentication.

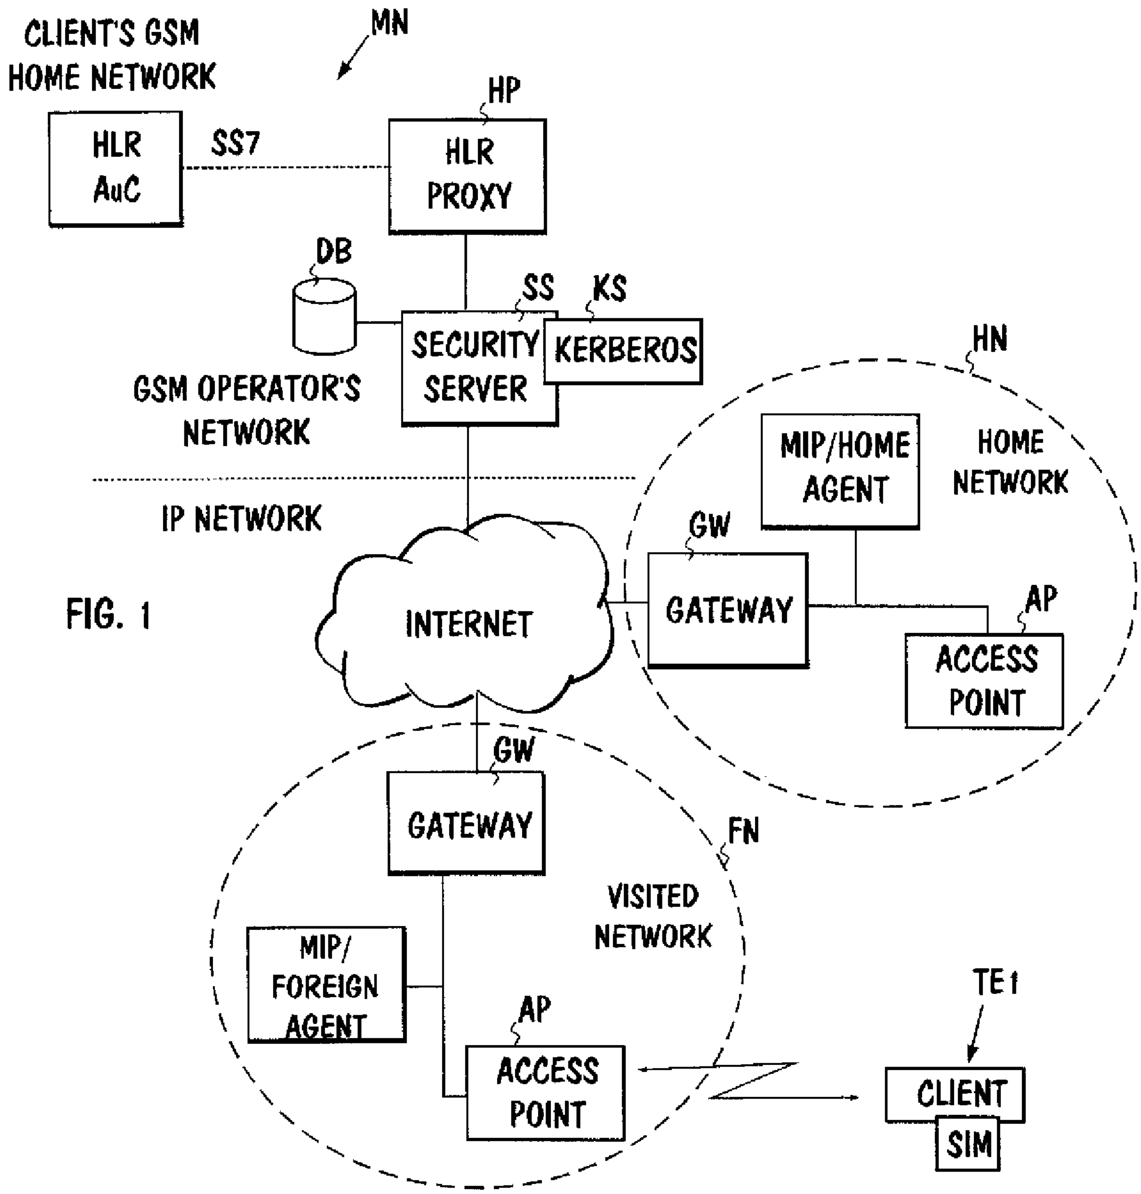
CLIENT'S GSM HOME NETWORK
MN
HP
HLR AuC
SS7
HLR PROXY
DB
SS
KS
SECURITY SERVER
KERBEROS
GSM OPERATOR'S NETWORK
HN
MIP/HOME AGENT
HOME NETWORK
IP NETWORK
GW
INTERNET
GATEWAY
AP
ACCESS POINT
GW
GATEWAY
FN
VISITED NETWORK
MIP/ FOREIGN AGENT
AP
TE1
ACCESS POINT
CLIENT
SIM

FIG. 1

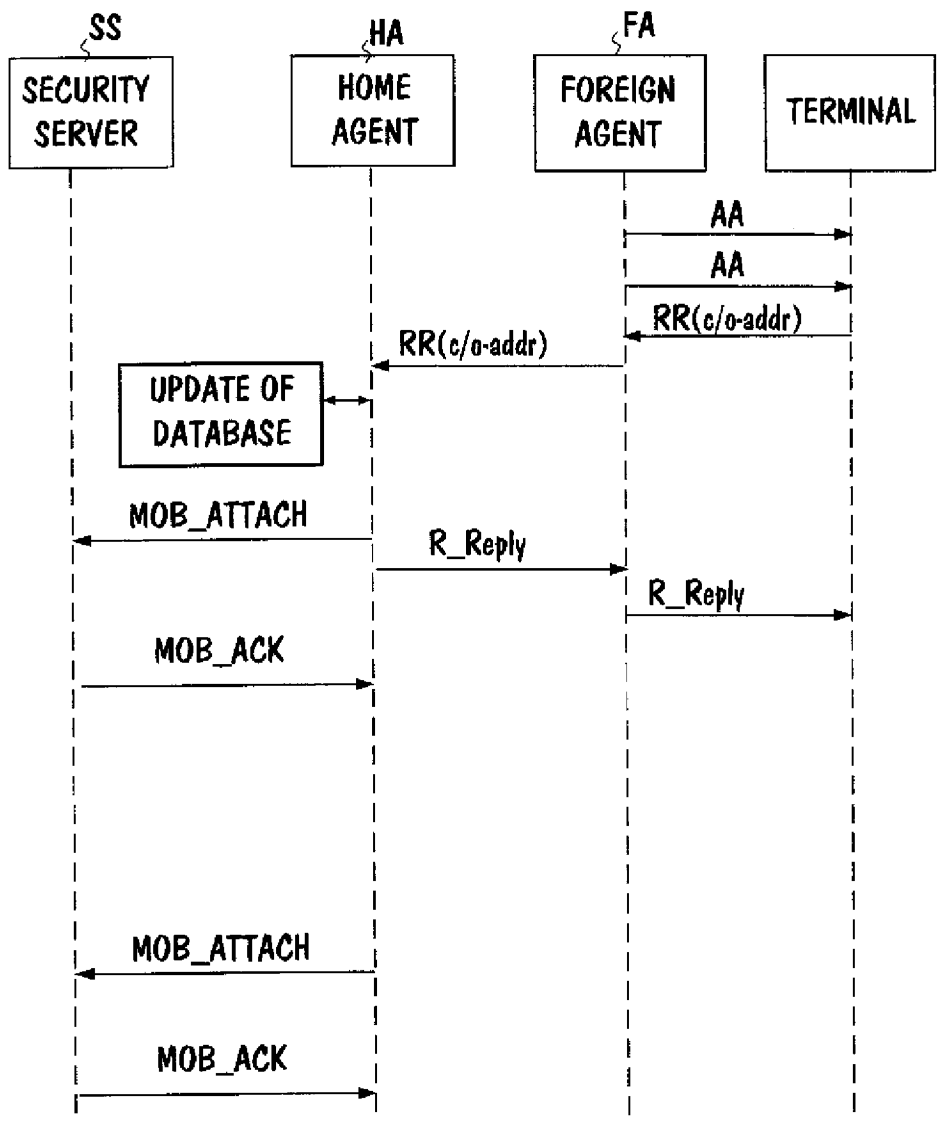
SS
HA
FA
SECURITY SERVER
HOME AGENT
FOREIGN AGENT
TERMINAL
AA
AA
RR(c/o-addr)
RR(c/o-addr)
UPDATE OF DATABASE
MOB_ATTACH
R_Reply
R_Reply
MOB_ACK
MOB_ATTACH
MOB_ACK

FIG. 2

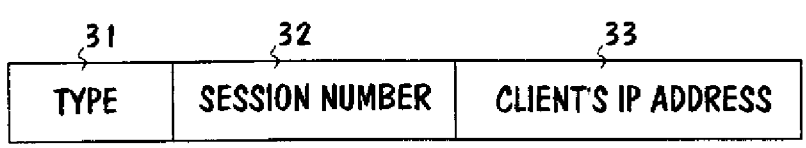
31
32
33
TYPE
SESSION NUMBER
CLIENT'S IP ADDRESS

FIG. 3

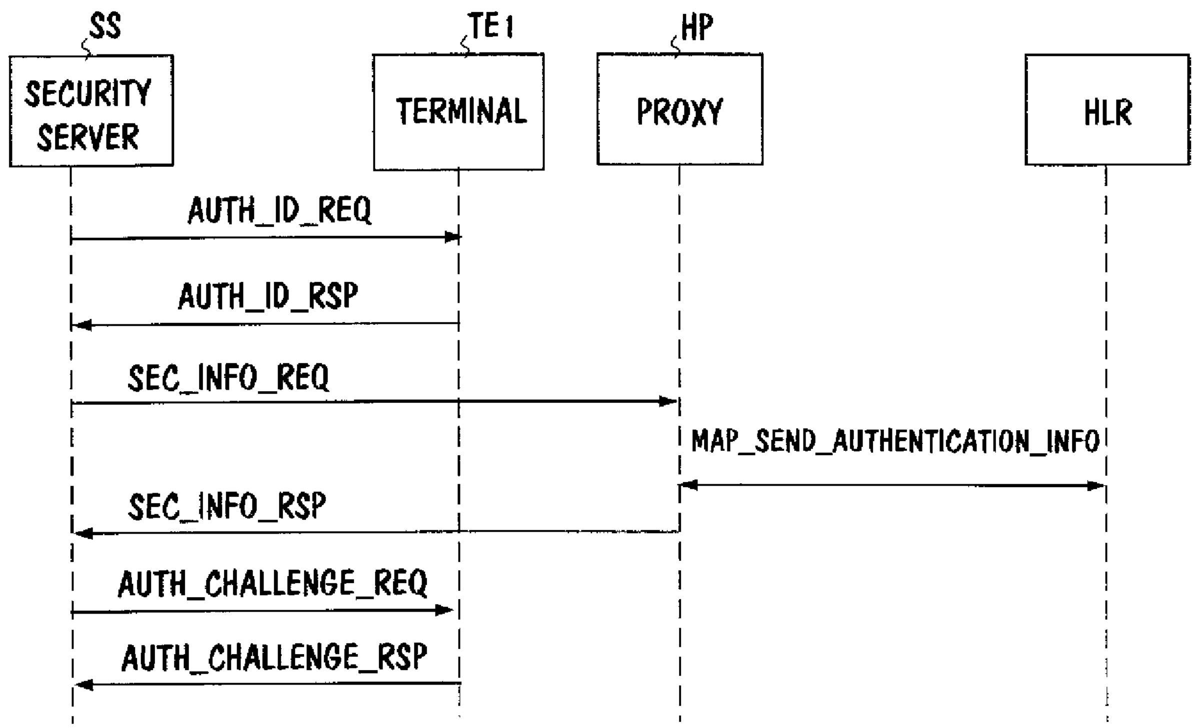

FIG. 4

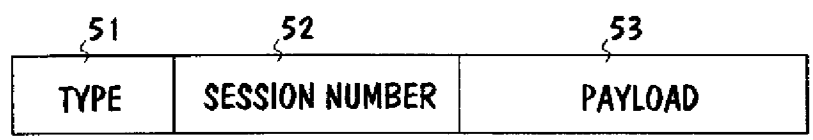

FIG. 5

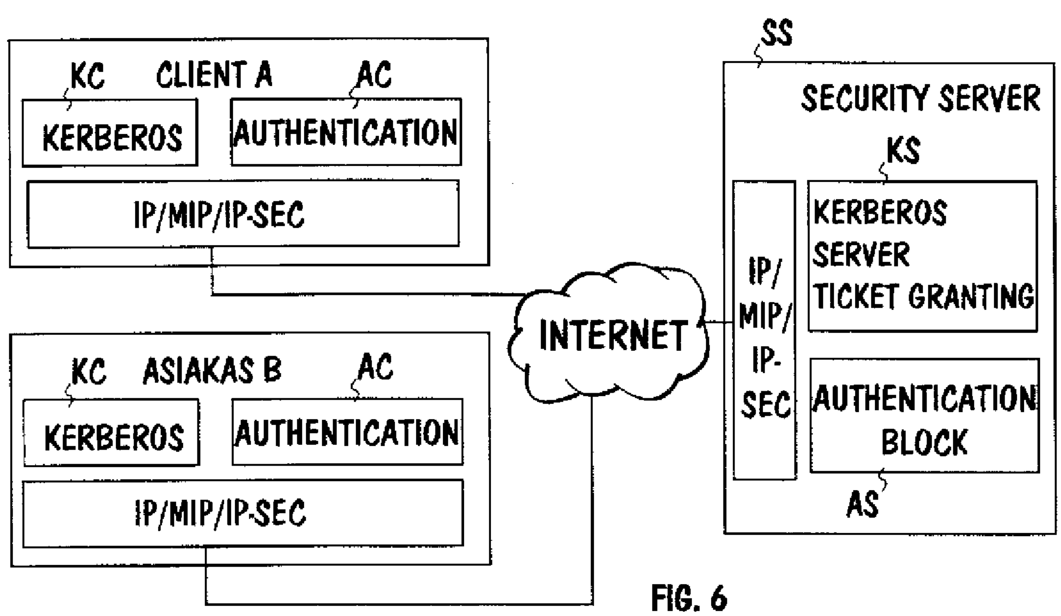

FIG. 6

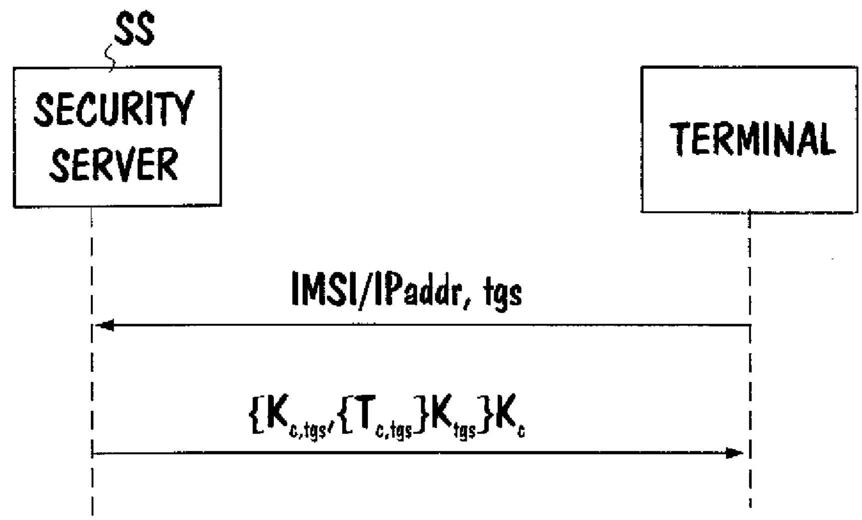

FIG. 7

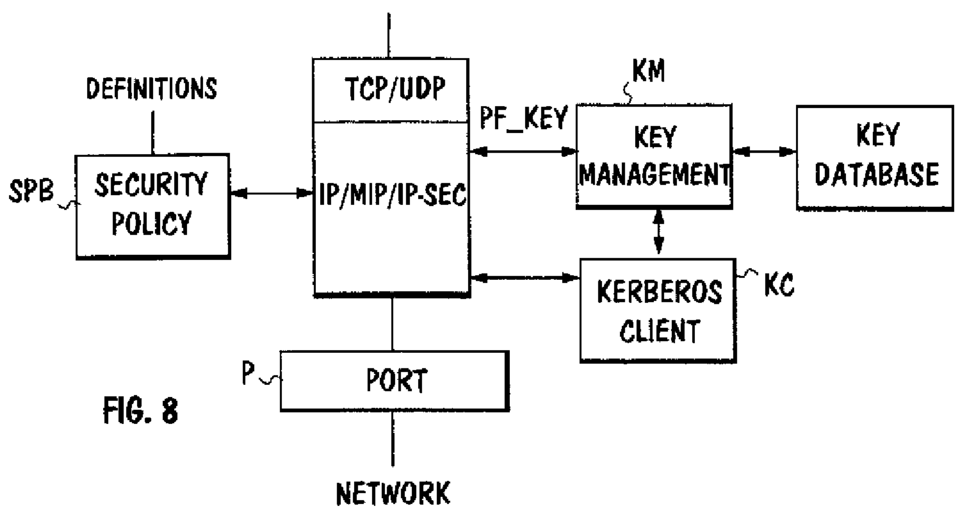

FIG. 8

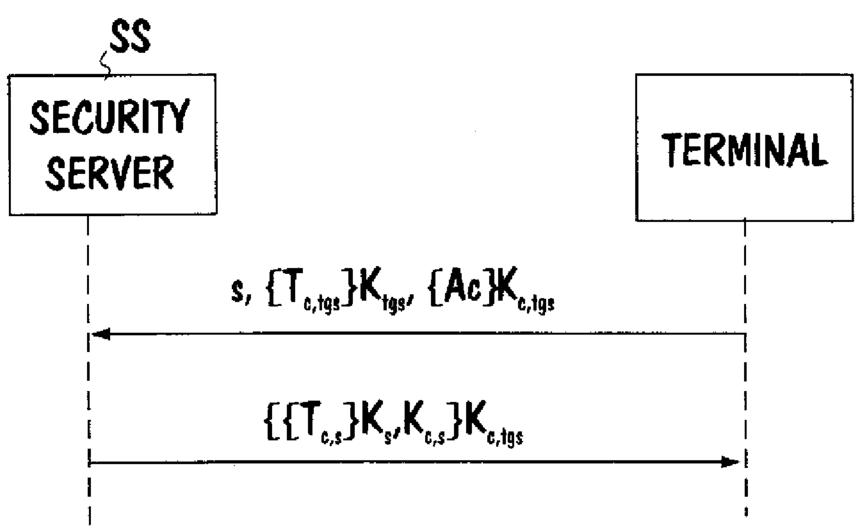

FIG. 9

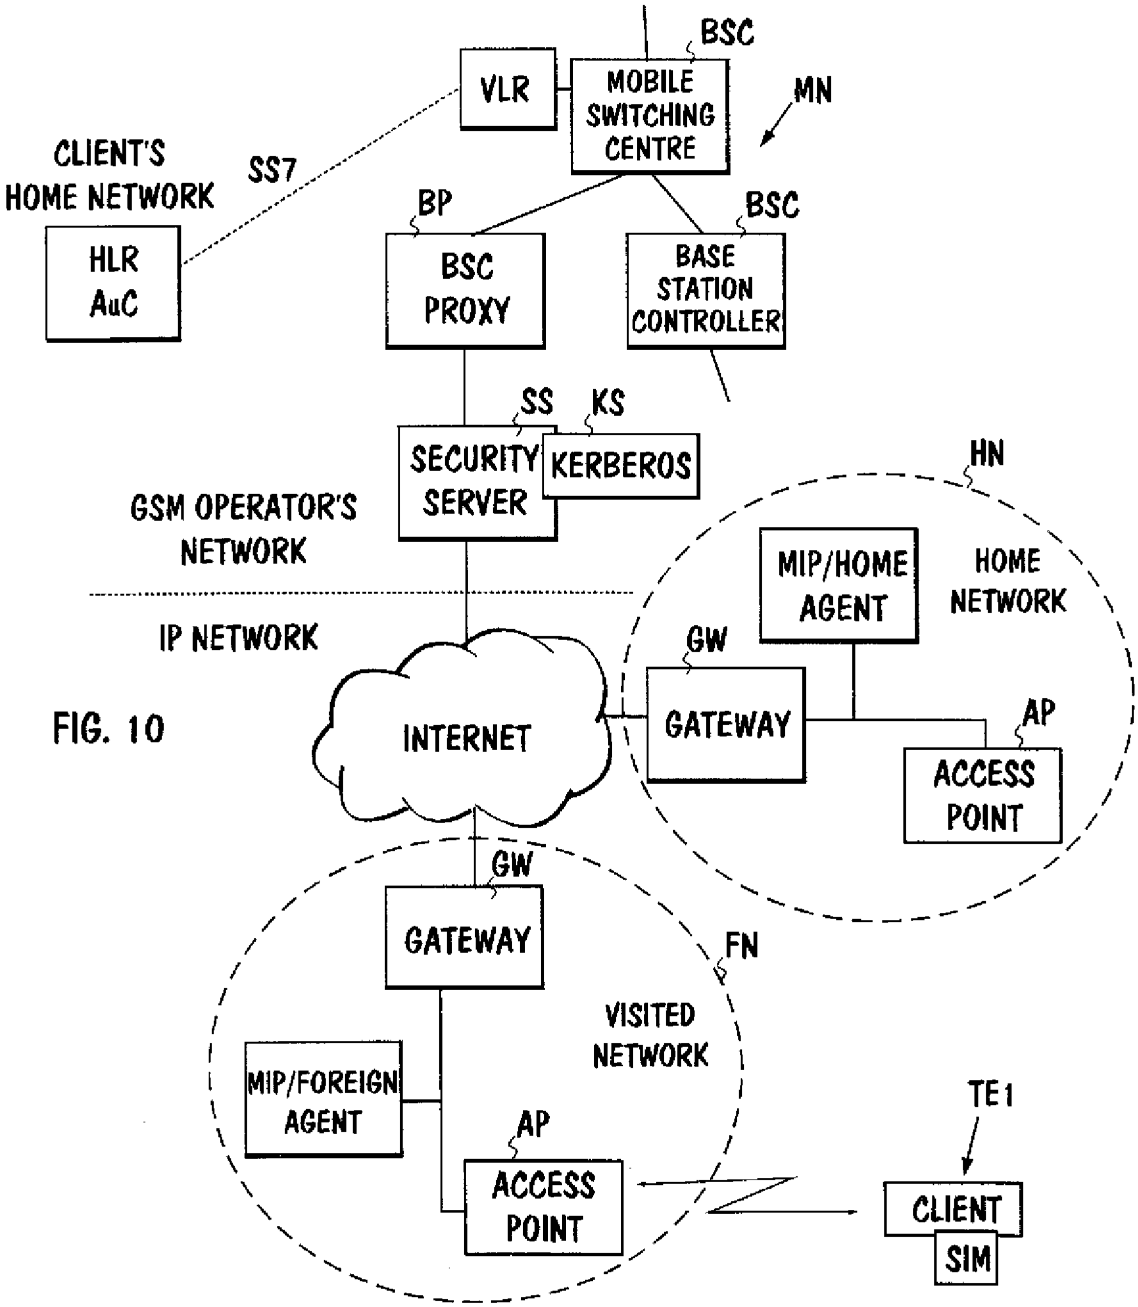
BSC
VLR
MOBILE SWITCHING CENTRE
MN
CLIENT'S HOME NETWORK
SS7
BP
BSC
HLR AuC
BSC PROXY
BASE STATION CONTROLLER
SS
KS
SECURITY SERVER
KERBEROS
HN
GSM OPERATOR'S NETWORK
MIP/HOME AGENT
HOME NETWORK
IP NETWORK
GW
INTERNET
GATEWAY
AP
ACCESS POINT
GW
GATEWAY
FN
VISITED NETWORK
MIP/FOREIGN AGENT
AP
ACCESS POINT
TE1
CLIENT
SIM

FIG. 10

PCT

WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 7 : H04Q 7/38

A3

(11) International Publication Number: WO 00/02406

(43) International Publication Date: 13 January 2000 (13.01.00)

(21) International Application Number: PCT/FI99/00565

(22) International Filing Date: 24 June 1999 (24.06.99)

(30) Priority Data:
981565 7 July 1998 (07.07.98) FI

(71) Applicant (*for all designated States except US*): NOKIA NETWORKS OY [FI/FI]; Keilalahdentie 4, FIN–02150 Espoo (FI).

(72) Inventor; and
(75) Inventor/Applicant (*for US only*): EKBERG, Jan–Erik [FI/FI]; Seljatie 1 A 5, FIN–00320 Helsinki (FI).

(74) Agent: PATENT AGENCY COMPATENT LTD.; Pitkänsillaranta 3 B, FIN–00530 Helsinki (FI).

(81) Designated States: AE, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, US, UZ, VN, YU, ZA, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SL, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

Published
With international search report.
Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.
In English translation (filed in Finnish).

(88) Date of publication of the international search report: 24 February 2000 (24.02.00)

(54) Title: SYSTEM AND METHOD FOR AUTHENTICATION IN A MOBILE COMMUNICATIONS SYSTEM

(57) Abstract

The invention concerns authentication to be performed in a telecommunications network, especially in an IP network. To allow a simple and smooth authentication of users of IP networks in a geographically large area, the IP network's terminal (TE1) uses a subscriber identity module (SIM) as used in a separate mobile communications system (MN), whereby a response may be determined from the challenge given to the identity module as input. The IP network also includes a special security server (SS), to which a message about a new user is transmitted when a subscriber attaches to the IP network. The subscriber's authentication information containing at least a challenge and a response is fetched from the said mobile communications system to the IP network and authentication is carried out based on the authentication information obtained from the mobile communications system by transmitting the said challenge through the IP network to the terminal, by generating a response from the challenge in the terminal's identity module and by comparing the response with the response received from the mobile communications system. Such a database (DB) may also be used in the system, wherein subscriber–specific authentication information is stored in advance, whereby the information in question need not be fetched from the mobile communications system when a subscriber attaches to the network.

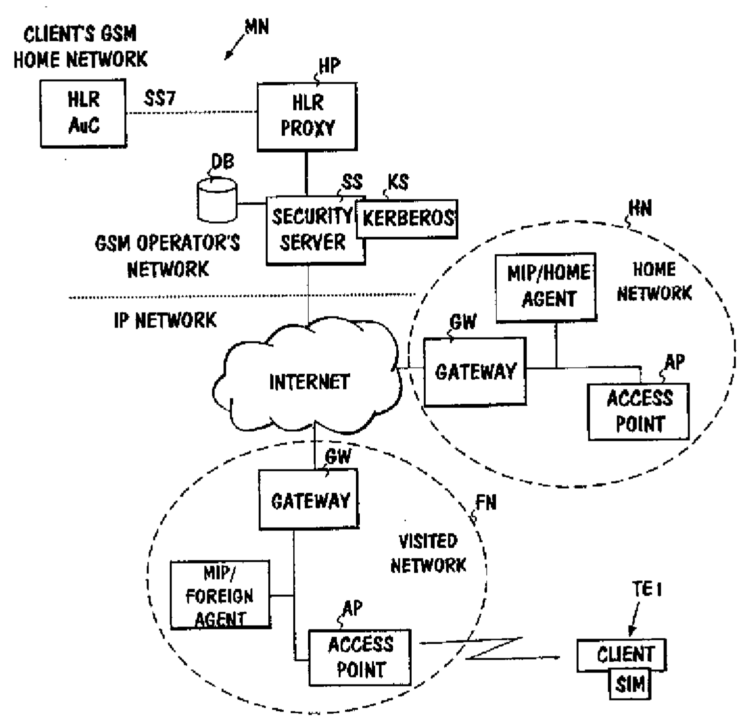

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece			TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	NZ	New Zealand		
CM	Cameroon			PL	Poland		
CN	China	KR	Republic of Korea	PT	Portugal		
CU	Cuba	KZ	Kazakstan	RO	Romania		
CZ	Czech Republic	LC	Saint Lucia	RU	Russian Federation		
DE	Germany	LI	Liechtenstein	SD	Sudan		
DK	Denmark	LK	Sri Lanka	SE	Sweden		
EE	Estonia	LR	Liberia	SG	Singapore		

INTERNATIONAL SEARCH REPORT

International application No.
PCT/FI 99/00565

A. CLASSIFICATION OF SUBJECT MATTER

IPC7: H04Q 7/38
According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC7: H04Q

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

SE,DK,FI,NO classes as above

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 5729537 A (LARS AXEL BILLSTÖM), 17 March 1998 (17.03.98) --	1-22
A	WO 9745814 A1 (VAZVAN, BEHRUZ), 4 December 1997 (04.12.97) --	1-22
P,A	WO 9844402 A1 (BRITISH TELECOMMUNICATIONS PUBLIC LIMITED COMPANY), 8 October 1998 (08.10.98) --	1-22
P,A	WO 9832301 A1 (TELEFONAKTIEBOLAGET LM ERICSSON (PUBL)), 23 July 1998 (23.07.98) --	1-22

[X] Further documents are listed in the continuation of Box C.   [X] See patent family annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" erlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance: the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance: the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
16 December 1999	03 -01- 2000
Name and mailing address of the ISA/ Swedish Patent Office Box 5055, S-102 42 STOCKHOLM Facsimile No. +46 8 666 02 86	Authorized officer Benny Andersson/mj Telephone No. +46 8 782 25 00

INTERNATIONAL SEARCH REPORT

International application No.
PCT/FI 99/00565

C (Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
P,A	US 5864757 A (JOHN PATRICK PARKER), 26 January 1999 (26.01.99)	1-22

--

INTERNATIONAL SEARCH REPORT

Information on patent family members

02/12/99

International application No.
PCT/FI 99/00565

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 5729537 A	17/03/98	AU 3199697 A	07/01/98
		CA 2258036 A	18/12/97
		CN 1227688 A	01/09/99
		EP 0904643 A	31/03/99
		WO 9748208 A	18/12/97
WO 9745814 A1	04/12/97	FI 962553 A	25/11/97
		FI 970767 A	20/10/97
		FI 971009 A	26/04/97
		FI 971248 A	26/04/97
WO 9844402 A1	08/10/98	AU 6414098 A	22/10/98
		GB 9800808 D	00/00/00
WO 9832301 A1	23/07/98	AU 5684698 A	07/08/98
		EP 0953265 A	03/11/99
US 5864757 A	26/01/99	AU 1409997 A	03/07/97
		CA 2239550 A	19/06/97
		EP 0867099 A	30/09/98
		IL 124872 D	00/00/00
		JP 11501182 T	26/01/99
		WO 9722221 A	19/06/97

PCT

WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 7 : H04L 29/06, 12/18, H04N 7/15

A1

(11) International Publication Number: WO 00/36804

(43) International Publication Date: 22 June 2000 (22.06.00)

(21) International Application Number: PCT/GB99/03871

(22) International Filing Date: 19 November 1999 (19.11.99)

(30) Priority Data:
9826158.9 27 November 1998 (27.11.98) GB

(81) Designated States: AU, CA, JP, SG, US, European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).

Published
With international search report.
Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.

(71) Applicant *(for all designated States except US)*: BRITISH TELECOMMUNICATIONS PUBLIC LIMITED COMPANY [GB/GB]; 81 Newgate Street, London EC1A 7AJ (GB).

(72) Inventors; and
(75) Inventors/Applicants *(for US only)*: BELL, Sarah [GB/GB]; 9 Brenda Gautrey Way, Cottenham, Cambridge CB4 8XW (GB). ING, Sarom [GB/GB]; 112 Spring Road, Ipswich, Suffolk IP4 2RR (GB). RUDKIN, Steven [GB/GB]; 52 Corder Road, Ipswich, Suffolk IP4 2XD (GB).

(74) Agent: SHELLEY, Mark, Raymond; BT Group Legal Services, Intellectual Property Dept., Holborn Centre, 8th floor, 120 Holborn, London EC1N 2TE (GB).

(54) Title: ANNOUNCED SESSION DESCRIPTION

(57) Abstract

The invention provides a method of announcing a description of a media session, for example a multimedia conference. In one respect, the invention provides a modular method of announcing media sessions. This method comprises the steps of generating a first base module (410) having a first data structure comprising user oriented data relevant to the media session; generating at least one media module (421, 422, 423) having a second data structure comprising media oriented data necessary for a user to receive a respective media stream of the media session; providing a link between the first base module and the at least one media module; and, announcing the media session by making at least the first base module available to potential recipients of the media session, wherein the link between the first base module and the at lest one media module permits a user to access the at least one media module and subsequently receive the media stream.

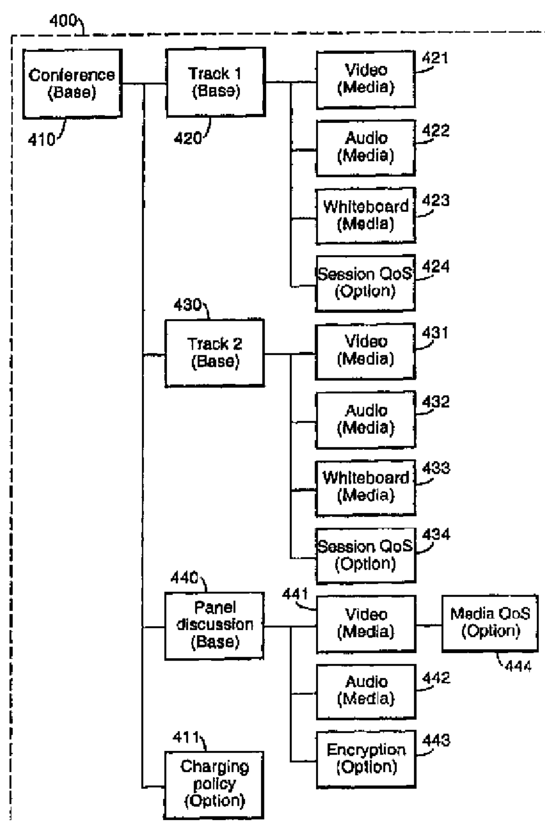

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	**ES**	Spain	**LS**	Lesotho	**SI**	Slovenia
AM	Armenia	**FI**	Finland	**LT**	Lithuania	**SK**	Slovakia
AT	Austria	**FR**	France	**LU**	Luxembourg	**SN**	Senegal
AU	Australia	**GA**	Gabon	**LV**	Latvia	**SZ**	Swaziland
AZ	Azerbaijan	**GB**	United Kingdom	**MC**	Monaco	**TD**	Chad
BA	Bosnia and Herzegovina	**GE**	Georgia	**MD**	Republic of Moldova	**TG**	Togo
BB	Barbados	**GH**	Ghana	**MG**	Madagascar	**TJ**	Tajikistan
BE	Belgium	**GN**	Guinea	**MK**	The former Yugoslav Republic of Macedonia	**TM**	Turkmenistan
BF	Burkina Faso	**GR**	Greece			**TR**	Turkey
BG	Bulgaria	**HU**	Hungary	**ML**	Mali	**TT**	Trinidad and Tobago
BJ	Benin	**IE**	Ireland	**MN**	Mongolia	**UA**	Ukraine
BR	Brazil	**IL**	Israel	**MR**	Mauritania	**UG**	Uganda
BY	Belarus	**IS**	Iceland	**MW**	Malawi	**US**	United States of America
CA	Canada	**IT**	Italy	**MX**	Mexico	**UZ**	Uzbekistan
CF	Central African Republic	**JP**	Japan	**NE**	Niger	**VN**	Viet Nam
CG	Congo	**KE**	Kenya	**NL**	Netherlands	**YU**	Yugoslavia
CH	Switzerland	**KG**	Kyrgyzstan	**NO**	Norway	**ZW**	Zimbabwe
CI	Côte d'Ivoire	**KP**	Democratic People's Republic of Korea	**NZ**	New Zealand		
CM	Cameroon			**PL**	Poland		
CN	China	**KR**	Republic of Korea	**PT**	Portugal		
CU	Cuba	**KZ**	Kazakstan	**RO**	Romania		
CZ	Czech Republic	**LC**	Saint Lucia	**RU**	Russian Federation		
DE	Germany	**LI**	Liechtenstein	**SD**	Sudan		
DK	Denmark	**LK**	Sri Lanka	**SE**	Sweden		
EE	Estonia	**LR**	Liberia	**SG**	Singapore		

ANNOUNCED SESSION DESCRIPTION

The present invention relates to the announcement of media stream connections for a media session over a communications network.

Multicast transmissions are becoming increasingly common on the Internet. In contrast to standard Internet Protocol (IP) point to point transmissions (unicast), IP multicast allows the simultaneous transmission of information to a group of recipients from a single source. Routing support for IP multicast transmissions is provided by the MBone (IP Multicast Backbone) which is a virtual network layered on top of the Internet.

IP multicast allows real-time communications over wide area IP networks and typical transmissions include video and audio conferencing, live multimedia training, university lectures and transmission of live television programmes.

A multicast transmission usually consists of a multimedia session made up of several individual media streams typically carrying video, audio, whiteboard or raw data. Some sessions are persistent, but the majority exist for a specific period of time, although need not be continuous. Multicast based transmissions on the MBone differ from unicast IP transmissions in that any user receiving the transmission can join the session (unless the transmission is encrypted) and to receive a transmission, a user need only know the appropriate transmission address and timing information.

Prior to a multicast transmission an appropriate announcement containing a session description is made, usually at an IP group multi-cast address. Standard session descriptions are generated using a Session Description Protocol (SDP), as defined in the Internet Engineering Task Force's draft RFC 2327. SDP is a simple ASCII text based protocol that is used to describe real time multimedia sessions and their related scheduling information. SDP messages are wrapped in a carrier protocol, known as a Session Announcement Protocol (SAP), which, in addition to containing the necessary

IP addressing and routing information for transmission across the Internet or MBone, allows the SDP message to be encrypted, signed or compressed. An announcement can then be sent at regular intervals to the announcement group address. As an alternative to SAP, a session may be announced by placing an SDP message on a World Wide Web site (WWW) or by sending it to individuals by email or as a unicast transmission inviting them to participate.

An SDP message conveys information about each media stream in the multicast multimedia session to allow the recipients to participate in the session. A typical SDP message will include the session name and purpose, the time(s) and date(s) the session will be active, the component media streams of the session and information required to participate in each media stream (IP multicast address, port, media format). The SDP message may also include details of the session's bandwidth requirements, an encryption key necessary to participate in a secure multicast transmission using public key encryption, contact information for the organiser of the multicast session, and a Unique Resource Indicator (URI) pointing to a WWW or an Intranet web site where further information on the session may be found, for example, background information relating to the conference.

The level of participation a user may make in a session or stream depends on its purpose. In a multicast television session, typically users would only be able to receive the session streams whilst in a multicast conference session the communication would be bi-directional with a central server (such as group address 120) receiving each participants transmissions and relaying them to the other participants. The level of participation expected of a user in a session or stream may be explicitly stated in the session description or it may be inherent from the session description, for example when a receive-only application is associated with a media stream type in the session description.

A common front end interface used by multicast end users is known as Session Directory Rendezvous (SDR). This interface takes the received announcements, decodes the SDP message and displays the names of those sessions that are still current in a list. The end user may then select one of the listed announcements to view further technical and user-oriented details of the announced session. From the displayed information, the end user can then select to join individual streams of the session or to join the entire session. Once the streams to be joined are selected, SDR starts the necessary multicast-enabled multimedia application on the end user's computer, such as Vic and Vat, and passes the relevant stream information (a transport port address) from the announcement onto the application allowing the application to establish the link to the associated IP multicast address and participate in the stream at transmission time. Having initiated the applications and passed the relevant transport port address SDR plays no further part in the session.

Recent increased usage and demand for (multi)media sessions has highlighted a number of limitations in SDP. SDP limits session descriptions to defining a session having a single set of timings that apply to all of the streams within it. A session in which a stream starts mid-way through the transmission cannot easily be described using SDP. The structure of a session description written in SDP must be a simple linear list of streams which may not reflect the intuitive structure of a complex session. SDP supports a limited and predefined set of applications which can receive the streams and a limited and predefined set of transport mechanisms (e.g. Simple layering, RTP and UDP). As guaranteed Quality of Service (QoS) is becoming more and more desirable to the consumer and the supplier, the need to define QoS policies for the entire session and individual streams in terms of required system resources, bandwidth requirements and supported applications also needs to be met. There may also be requirements on the prioritisation of streams and subsessions or more complicated rules about receiving

streams. A further requirement on the part of the supplier will be the need for charging facilities permitting the charging of an end user for a multicast transmission to which they subscribe according to the QoS and types of streams received etc. There is little scope to include information about QoS policies or charging within the conventional structure of an SDP session description, or any metadata about the session.

According to a first aspect of the present invention there is provided, a method of announcing a description of a media session, comprising the steps of:

generating a first base module having a first data structure comprising user oriented data relevant to the media session;

generating at least one media module having a second data structure comprising media oriented data necessary for a user to receive a respective media stream of the media session;

providing a link between the first base module and the at least one media module; and,

announcing the media session by making at least the first base module available to potential recipients of the media session,

wherein the link between the first base module and the at least one media module permits a user to access the at least one media module and subsequently receive the media stream.

The present invention provides a modular description system for a media session in which session descriptions are constructed in a hierarchical manner providing a plurality of levels of information concerning the constituent parts of the described session.

A problem faced with the current distribution of announcements from the single announcement group address is that there is a limit to the size of each announcement and the frequency with which each can be sent out. In the present invention, it is

possible to provide a modular description system in which a distributed announcement contains links available to the end user to other portions of the announcement which have not been transmitted.

Preferably, the method further comprises the steps of: generating a second base module, the second base module containing user orientated data relating to a sub-session of the media session; linking the second base module to the first base module; and, linking said at least one media module to the second base module.

In preferred embodiments, the method further comprises the steps of: generating at least one options module having a third data structure comprising data relating to service level criteria required to participate in the media session; and, linking the or each options module to a respective base module.

The data contained in the options module may relate to a quality of service policy to be used by the media session or a part thereof. Alternatively, the data contained in the options module may relate to a security system to be used by the media session or a part thereof. The data contained in the options module may further relate to a charging system to be used by the media session or a part thereof.

In preferred embodiments, one or more media module(s) comprise data necessary for a user to receive a layered media stream of a respective media session; and said method further comprises the step of linking the or each media module to one or more respective options module(s) containing data relating to a layered mechanism of the respective layered media stream necessary for a party to participate in the layered media stream.

The media session may be announced by transmitting all of the constituent modules of the session description. Alternatively, the media session may be announced by transmitting only some of the constituent modules of the session description, with the

remaining modules of the session description being subsequently accessible by a user using one or more links provided in the modules transmitted. The remaining modules of the session description may be held on one or more servers and the one or more links to the remaining modules are in the form of URI pointers. Modules of the session description contain links to modules which are generated subsequent to the announcement.

According to a second aspect of the invention there is provided a computer readable storage medium containing data defining at least a part of a description of a media session, the session description comprising:-

a first base module having a first data structure comprising user oriented data relevant to the media session;

at least one media module having a second data structure comprising media oriented data necessary for a user to receive a respective media stream of the media session;

a link between the first base module and the at least one media module;

Another problem faced by providers of current (multi)media sessions and the developers of the associated (multi)media applications is the spread of skills required to implement an application that can initiate and manage a real-time data connection over a communications network and perform the (multi)media functions the end user would expect. For example, developers of multimedia applications require teams with skills in audio and video coding, network transport protocols, real time programming, user interface design and integration techniques. The session description of the present invention simplifies this process by allowing the necessary communication channels and media streams to be identified in the session description. This information is used by generic middleware in the form of a session control and communications manager to dynamically instantiate the respective streams and channels for the applications at run

time.

Furthermore, until now the only way a QoS policy could be implemented was to process a session description to determine which streams of a session could or should be run and then to initialise the applications so they connect to the respective streams. This required the communications manager not only to know about the session requirements and available system resources but also the capabilities of each application.

In a preferred example of the present invention the media modules of a session description are checked by the respective multimedia client application prior to QoS management, thereby reducing the workload of the communications manager, that is to say the respective client applications determine whether the media modules can be supported. Furthermore, applications need only request streams from the session control system associated with the client since the session control now handles centrally the creation and management of streams in real time. This aspect is also the subject of our co-pending UK patent application 9826157.1.

The present invention simplifies application development and service provision. A further problem is that applications should be able to adapt to available network and host resources. This is particularly important for multi-party applications operating in heterogeneous environments where each party may have different resources available to them. Furthermore the nature of the heterogeneity may vary over the lifetime of the session, for example as network congestion varies or as the terminal resources are shared with other applications or other users. The present invention is able to use a QoS policy incorporated within the session description to prioritise the allocation of resources and to determine whether participation in the session is viable.

A still further problem is that the application developer and service provider typically need to address security and charging requirements. The present invention

allows security and charging policies to be incorporated within the session description for use within the session control system to invoke appropriate charging and security procedures. Instead of having to develop security and charging functions the application developer and service provider need only specify appropriate policies.

In the present invention application development is simplified by using the session description to drive the dynamic management of communication channels and to adapt to available resources. It also reduces the problem of handling charging and security requirements to a matter of specifying charging and security policies within the session description.

An example of the present invention will now be described in detail with reference to the accompanying drawings, in which:

Figure 1 is a schematic diagram illustrating a multicast transmission across the MBone;

Figure 2 is a schematic diagram illustrating the distribution of an SDP announcement;

Figure 3 is a block diagram of a modular session description of a simple session generated in accordance with the present invention;

Figure 4 is a block diagram of a modular session description of a complex session generated in accordance with the present invention;

Figure 5 is a schematic diagram of a system for managing media stream connections;

Figure 6 is a flow chart illustrating the steps involved in managing a media session according to the system of Figure 5; and,

Figure 7 is a flow chart further illustrating a parsing step of Figure 6.

An example of an IP multicast transmission system is described with reference to Figure 1. Prior to a multicast transmission, an appropriate announcement containing

a session description is made, thereby allowing end users 110a-110e to elect to receive the transmission. Each end user electing to receive the transmission is linked to a group IP Multicast address 120 associated with the transmission. At the transmission time of the multicast session, the session streams are transmitted from a source 130, or a plurality of sources, to the group address. At the group address, the transmission is disseminated along the links 140 to those end users who have elected to receive it (in this example end users 110a-110c).

An example of an announcement and election system is described with reference to Figure 2. Most public multicast sessions are announced at a single group IP multicast address 200 dedicated to the transmission of announcements to multicast end users. End users 210a-210e electing to receive the announcements are linked to the announcement group address and, in the same way as an actual session transmission, each announcement arriving at the announcement group address is disseminated to the end users. A front end interface 220 on each end user's computer displays information obtained from the associated session description for each announcement. The minimum information a session description may contain is a time and date that the session will be active and the group IP multicast address(es) from which the end user may elect to receive one or more media streams and to which they could send their own streams for the session. Using the front end interface, an end user can select the announced session(s), or their component stream(s) they wish to participate in.

Figure 3 is a block diagram of a session description 300 for a simple multicast television session. The session description 300 comprises a base module 310 linked to a media module 320.

The base module 310 contains user oriented data relating to the session including the title and timing information. The base module 310 may also include a

description or abstract, contact information about the organiser and a WWW or an intranet URI pointing to a web site containing further information. Ideally, the base module 310 should contain enough information for the user to decide if they are interested in participating in the session.

The media module 320 contains announcement data relating to a video stream of the session. The media module 320 contains the technical information (data) necessary for the user to receive the associated media stream. In particular, connection, timing and media format details are provided.

A first example of a session description 300 generated for transmission to end users is shown below:

```
(
type=(base)
id=(310)
info=(title="live multicast television session")
source=(name="A.Sender" email=asender@tx.com)
media=(video=(client=odbits0.16))
time=(length=50m repeat=continuous)
category=("Entertainment")
options=(none)
modules=(m=320)
)

(
type=(media)
id=(320 310)
media=(video=(client=odbits0.16))
connection=(229.1.1.2/7000)
time=(length=50m)
)
```

Session description example 1

The base module 310 has a unique identifier (id field) used in the generation of links between two modules during the processing of the session description. The modules field of the base module 310 lists the type and unique identifier of the media

module 320 linked to the base module 310. The second identifier in the id field of the media module 320 is the unique identifier belonging to the base module 310 linking the media module back to the base module 310. By extension, these two-way links permit a module tree to be traversed from a base module downwards or from a media module upwards. The use of this feature is described later with reference to session description example 4.

The connection field of the media module 320 contains the IP multicast address and port number from which the media stream can be received.

Figure 4 is a block diagram of a session description 400 for a complex multicast session of a multimedia conference with two tracks, or sub-sessions, and a panel discussion. Each track provides multiparty video and audio conferencing and a shared whiteboard for leaving notes and messages. The panel discussion is encrypted and the whole conference is subject to a subscription fee payable in advance by each participant.

The session description 400 contains a top level base module 410 linked to further base modules 420, 430, 440 and an options module 411. The top level base module 410 contains data relating to the overall session including its name, purpose and timing information. The options module 411 contains details of the payment mechanism for subscription fees.

Each further base module 420, 430, 440 relates to a subsession of the conference. Base module 420 relates to the first track of the conference. The base module 420 is linked to media modules 421-423, each containing connection, timing and media format data for respective video, audio and whiteboard streams.

The base module 420 is also linked to options module 424 which contains data relating to a QoS policy for the first track defining which media modules are optional and which are mandatory for a participant of the first track. The mandatory list contains

identifiers of those media modules which are needed for the session or subsession to operate correctly whilst the optional list contains identifiers of the media modules that are not necessary for the session or subsession to operate correctly if system resources are scarce.

The base module 430 relates to the second track of the conference. It is linked to media modules 431-433, each containing connection, timing and media format details for respective video, audio and whiteboard streams. The base module 430 is also linked to options module 434 which contains data relating to a QoS policy for the second track defining which media modules are optional and which are mandatory for a participant of the second track. Base module 440 relates to the panel discussion. It is linked to media modules 441 and 442, each containing connection, timing and media format details for respective video and audio streams of the panel discussion. The base module 440 is also linked to options module 443 which contains encryption details (ie. how and where to get the necessary cryptographic keys) necessary for a participant to decode the panel discussion media streams 441, 442 according to a known encryption mechanism such as DES or public key encryption.

The video media stream defined in media module 441 is layered. Layering of media streams allows users with different system resources to receive as much of the stream as their system resources allows. Every user must receive the bottom layer of the stream containing the minimum stream data. However, if a user has sufficient free system resources they can receive the next layer up containing enhancements to the previous layer. Successive layers can be received enhancing the received media stream until the maximum number of layers is received or all free system resources capacity is used. The media module 441 is linked to an options module 444 which contains data on the layering necessary for the end user to be able to receive the layered stream correctly.

The portion of the session description 400 generated for modules 410, 411, 420 and 440 for transmission to end users is shown below in session description example 2.

```
( # overall conference session
  type=(base)
  id=(410)
  info=(title="Multimedia98 Conference")
  source=(owner="Joe Bloggs" email=joe@nowhere.com)
  media=(video=(client=RealPlayerG2) whiteboard=(client=wb))
  time(start="09:00 GMT 25/12/98" stop="13:00 GMT 25/12/98")
  options=(oc=411)
  modules=(b=420 b=430 b=440 oc=411)
)

( # conference track 1
  type=(base)
  id=(420 410)
  info=(title="MM98 Systems and Applications Track")
  source=(owner="Joe Bloggs" email=joe@nowhere.com)
  media=(video=(client=RealPlayerG2) whiteboard=(client=wb))
  time(start="09:00 GMT 25/12/98" stop="11:00 GMT 25/12/98")
  options=(osq=424)
  modules=(m=421 m=422 m=423 osq=424)
)

( # session QoS for track 1
  type=(option-sQoS)
  id=(424 420)
  mandatory=(421 422)
  optional=(423)
)

( # conference panel discussion
  type=(base)
  id=(440 410)
  info=(title="MM98 Panel Discussion")
  source=(name="Joe Bloggs" email=joe@nowhere.com)
  media=(video=(client=RealPlayerG2) whiteboard=(client=wb))
  time(start="11:00 GMT 25/12/98" stop="13:00 GMT 25/12/98")
  options=(osec=443)
  modules=(m=441 m=442 osec=443)
)

( # video for panel discussion
  type=(media)
  id=(441 440)
  info=(title="MM98 Panel Discussion Video")
  source=(owner="Joe Bloggs" email=joe@nowhere.com)
  media=(video=(type=live client=RealPlayerG2))
  connection=(226.0.0.106/1010 policy=444)
```

```
    time=(start="11:00 GMT 25/12/98" stop="13:00 GMT 25/12/98")
)

( # media QoS policy for panel discussion video
  type=(option-mQoS)
  id=(444 440)
  mechanism=(layer=(base=226.0.0.106/1010 number=3))
)

( # encryption policy for panel discussion
  type=(option-sec)
  id=(443 440)
  participant=(member=w3c)
  publickey=(location=http://www.w3.org/members_only/)
  info=(location=http://www.w3.org/)
)

( # charging policy for entire conference
  type=(option-chg)
  id=(411 410)
  mechanism=(type=AAA)
  price=(fee=1000GBP)
  info=(location=http://www.aaa.net/)
)
```

Session description example 2

Where there is surplus network bandwidth available, complete session descriptions can be announced to end users who may then elect to receive the announced session or parts thereof. However, the individual modules of the session description do not need to be announced together. If the network bandwidth available for announcements restricts the size of session descriptions, only the top level base module may be announced. In this situation, the link between modules may be, for example, a URI to a WWW or an intranet web site or server, an email address, an IP multicast address, an FTP address or details of a file or database stored on a local computer system from which an interested user can obtain the remaining modules.

The following session description example illustrates how the above session description for base module 420 would be changed if media module 421 was stored on a WWW server:

```
( # conference track 1
  type=(base)
  id=(420 410)
  info=(title="MM98 Systems and Applications Track")
  source=(owner="Joe Bloggs" email=joe@nowhere.com)
  media=(video=(client=RealPlayerG2) whiteboard=(client=wb))
  time(start="09:00 GMT 25/12/98" stop="11:00 GMT 25/12/98")
  options=(osq=424)
  modules=(m=421 location=http://www.announce.org/cgi-bin/module.cgi?id=421
           m=421 m=423 osq=424)
)
```

Session description example 3

Furthermore, top level modules of a session description may be announced well in advance of the actual transmission, at a time where the final details of content are unknown, in which case the remaining levels may be made available from pre-announced links at a later time.

Figure 5 is a schematic diagram of a system for managing media stream connections at a terminal of an end user system according to the present invention.

The session control system 500 is linked to an announcement receiving interface 510 and one or more multicast-capable multimedia applications 520. The session control system 500 and the announcement receiving interface 510 are connected to a network interface 530 via which announcements may be received and multicast transmissions may be initiated and/or received.

Announcements received at the network interface 530 are routed to the receiving interface 510. The receiving interface 510 decodes each announcement to obtain the session description and displays the user oriented information from the one or more base modules in a list to the user. The user is able to select a session description from the list announcing a session they wish to receive. The selected description is passed to the session control system 500 which determines which of the user's multimedia applications 520 are required for participation in the described session, starts the applications and initiates and provides the necessary media streams

to the respective applications 520 via a communications manager 550.

The receiving interface 510 may be linked to other Internet communications applications 540 such as a WWW browser or an email client (not shown) which may be used to gather further information on the described session based on links provided in the session description. Also, where an incomplete set of base and/or media modules of a session description are received, the receiving interface 510 attempts to obtain the remaining modules using the Internet communications applications prior to passing it onto the session control system 500.

Figure 6 is a flow chart showing the steps taken by the session control system 500 upon receipt of a session description. The description is first parsed in step 600 to identify client applications for each media module. Once this is done a second parse is carried out where applications are launched in step 610, that is to say for each media module start the application specified in the client field if that application has not already been started. The portion of the session description relating to the respective media type, i.e. the media module, the base module directly above the media module, all other modules attached to that base module and any other options modules that apply, is passed to the corresponding application in step 620. Since the media modules are marked with appropriate client applications, each application will be able to select those media streams that it wants to participate in. The application replies to the session control system with a connection request specifying its requirements in the form of a list of identifiers of media modules from which streams are to be initiated in step 630. The connection request is assembled by the session control system in step 640 and the system then parses the session description to identify other applications to launch in step 645. If a further media type is found, steps 610 to 640 are repeated, otherwise the session control system uses the assembled connection requests to form a list of media modules. This list is passed, together with a session QoS policy, to the

communications manager, a system used in by the session control system, which determines according to the QoS policies and available system resources whether each connection request is viable.

The session QoS policy is constructed in two steps:- first, the multiple session QoS policies relevant for all the media modules to be initiated are combined into one session QoS policy: second, the resulting session QoS policy may be adapted to take account of (a) user default preferences (defined in a user profile), (b) a user's wish to determine the policy interactively, and (c) an application's default configuration (defined in the application profile(s)).

The communications manager responds to the session control system in step 650 with an indication of the viable media stream connection requests. If necessary, the session control system may contact a charging system to initiate accounting for the session prior to requesting the communications manager to create the viable media stream connections in step 660.

Once a session starts, each received data stream relating to the session is passed to the associated multimedia application in step 670 until the scheduled stream time ends in step 680 or the multimedia application requests to the session control system that the connection is terminated in step 690, at which point the session control system disconnects the connection in step 700.

Figure 7 is a flow chart showing the QoS management step 650 of Figure 6 in greater detail.

Having received the assembled list of connection requests, the communications manager matches each item of this list to a media profile in step 705. A media profile defines requirements which must be met for the requested media stream to operate on the end user's computer including the minimum network bandwidth needed for satisfactory reception of the stream.

A terminal profile is determined in step 710. The terminal profile defines the resources which are available at the end user's computer for use by the requested media streams. This includes available network bandwidth, free memory and disk space and available hardware such as monitor size, processor speed and free audio and video capture devices. The media profile of each connection request is compared against the available system resources defined by the terminal profile in step 720. If the terminal profile matches or exceeds the media profile, the connection request is declared viable in step 730 and the terminal profile is decremented accordingly for the remaining connection requests in step 740. Each connection request is processed until there are no remaining requests or until the media profile of a request exceeds the terminal profile. In this situation, the communications manager determines the optimum terminal profile the user's computer would have if all non-essential applications were not running in step 750 and whether the computer is capable of fulfilling the media profile in step 760. If the computer is capable of fulfilling the media profile, the communications manager attempts to free system resources from currently allocated streams or connection requests which have lower priority or by asking the user to terminate other non-essential applications running on the computer in step 770. Alternatively, this could be done by reducing the number of layers received from a layered stream transmission. If sufficient resources cannot be found an exception is reported to the user and the connection request is marked as unviable. If the media stream that cannot be received is defined as mandatory in a QoS policy for a media session or subsession, all the connection requests for that media session or subsession are cancelled in step 790. If, however, the media stream is optional, the communications manager continues processing further connection requests in step 720. Once all pending connection requests have been processed, the communications manager reports those that are viable to the session control system.

The processing of a session description will now be described with reference to Figure 4 and session description example 4 which is the session description generated for Track 1 (modules 410 and 420-424 of Figure 4).

```
( # overall conference session
  type=(base)
  id=(410)
  info=(title="Multimedia98 Conference")
  source=(owner="Joe Bloggs" email=joe@nowhere.com)
  media=(video=(client=RealPlayerG2) whiteboard=(client=wb))
  time(start="09:00 GMT 25/12/98" stop="13:00 GMT 25/12/98")
  options=(oc=0010)
  modules=(b=420 b=430 b=440 oc=411)
)

( # conference track 1
  type=(base)
  id=(420 410)
  info=(title="MM98 Systems and Applications Track")
  source=(owner="Joe Bloggs" email=joe@nowhere.com)
  media=(video=(client=RealPlayerG2) whiteboard=(client=wb))
  time(start="09:00 GMT 25/12/98" stop="11:00 GMT 25/12/98")
  options=(osq=424)
  modules=(m=421 m=422 m=423 osq=424)
)

( # video for track 1
  type=(media)
  id=(421 420)
  info=(title="MM98 Systems and Applications Track Video")
  source=(owner="Joe Bloggs" email=joe@nowhere.com)
  media=(video=(type=live client=RealPlayerG2))
  connection=(226.0.0.100/1000)
  time=(start="09:00 GMT 25/12/98" stop="11:00 GMT 25/12/98")
)

( # audio for track 1
  type=(media)
  id=(422 420)
  info=(title="MM98 Systems and Applications Track Audio")
  source=(owner="Joe Bloggs" email=joe@nowhere.com)
  media=(audio=(type=live format=g711))
  connection=(226.0.0.101/1001)
  time=(start="09:00 GMT 25/12/98" stop="11:00 GMT 25/12/98")
)

( # whiteboard for track 1
  type=(media)
  id=(423 420)
  info=(title="MM98 Systems and Applications Track Whiteboard")
  source=(owner="Joe Bloggs" email=joe@nowhere.com)
```

```
    media=(whiteboard=(client=wb))
    connection=(226.0.0.102/1002)
    time=(start="09:00 GMT 25/12/98" stop="11:00 GMT 25/12/98")
)

(  # session QoS for track 1
    type=(option-sQoS)
    id=(424 420)
    mandatory=(421 422)
    optional=(423)
)
```

Session description example 4

The session control system, having received the above session description, processes the tree structure of the session description starting at base module 410. The first module encountered is base module 420. As this is not a media module but it does have sub-modules, the session control system continues down this branch to media module.

The media field of the media module 421 already defines the multimedia client application required as RealPlayerG2 (a multimedia application of Real Networks Inc) thus the session control system ignores it and continues to the next media module. The media field of the media module 422 does not have a multimedia client application defined, however a format for the audio data is specified. The session control system recognises that this particular audio format can be supported by RealPlayerG2 so it amends the media field to read client=RealPlayerG2. The next media module 423 has already defined a client application as wb so it ignores this module, and it also ignores the option module 424.

The session control system parses the tree structure again in order to launch client applications. The first media module 421 specifies that RealPlayerG2 should be launched, hence the session control system launches the application on the end user's system and keeps a record of this activity. The second media module 422 specifies an

application that has already been launched and so the session control system ignores it and continues to the next media module. The media module 423 specifies that wb should be launched, so the session control system launches the application and keeps a record of this activity.

RealPlayerG2 is passed the media module 421, base module 420 and modules 422-424. The application processes the media modules given to determine which it can handle, and in this case it identifies 421 and 422. Having determined which streams it can handle, the application sends a connection request back to the session control system requesting connection to the media streams of modules 421 and 422. Similarly, wb is passed the media module 423, base module 420, modules 421-422, and the module 424. The application processes the given modules as described previously, and requests connection to the media stream of modules 423.

The above connection requests are assembled by the session control system into a list, this list is then passed to the communications manager along with the session QoS policy module 424. The communications manager determines whether each request is viable according to the steps of Figure 7.

Assuming there are sufficient resources for all the connection requests for mandatory media streams, the communications manager passes back a list of viable streams to the session control system which then processes the tree again to determine the connection data held in the connection field of each media module so it can request that the communications manager initiate a connection to the appropriate media stream for each of the viable connection requests according to the connection data. The session control system then manages the session and its media stream connections as is described with reference to steps 670 to 700 of Figure 6.

Due to the heterogeneity of the Internet and differing capabilities and operating environments of end user computer systems, the session control system described has

been implemented in Java (Java is a Trade Mark of Sun Microsystems Inc.). The announcement receiving interface, Session Directory, receives the announcements and passes those selected by the end user to the session control manager implemented as an application programming interface running as a background process on the end user's computer.

Whilst the present invention has been described with reference to the Internet and multicast transmissions, it will be apparent to the reader that the described modular session description and the session control system are applicable to the announcement and subsequent management of connections to media streams of a (multi)media session using other known transport mechanisms such as unicast.

Furthermore, although mechanisms for encryption, charging and other such services have not been explicitly described, it would be apparent to the reader that appropriate session descriptions and associated functions within the session control system for their processing could be readily implemented according to the mechanism required.

CLAIMS

1. A method of announcing a description of a media session, comprising the steps of:

generating a first base module having a first data structure comprising user oriented data relevant to the media session;

generating at least one media module having a second data structure comprising media oriented data necessary for a user to receive a respective media stream of the media session;

providing a link between the first base module and the at least one media module; and,

announcing the media session by making at least the first base module available to potential recipients of the media session,

wherein the link between the first base module and the at least one media module permits a user to access the at least one media module and subsequently receive the media stream.

2. A method according to claim 1, further comprising the steps of:

generating a second base module, the second base module containing user orientated data relating to a sub-session of the media session;

linking the second base module to the first base module; and,

linking said at least one media module to the second base module.

3. A method according to claim 1 or 2, further comprising the steps of:

generating at least one options module having a third data structure comprising data relating to service level criteria required to participate in the media session; and,

linking the or each options module to a respective base module.

4. A method according to claim 3, in which the data contained in the options module relates to a quality of service policy to be used by the media session or a part thereof.

5. A method according to claim 3 or 4, in which the data contained in the options module relates to a security system to be used by the media session or a part thereof.

6. A method according to any of claims 3 to 5, in which the data contained in the options module relates to a charging system to be used by the media session or a part thereof.

7. A method according to any preceding claim, wherein one or more media module(s) comprise data necessary for a user to receive a layered media stream of a respective media session; and said method further comprises the step of linking the or each media module to one or more respective options module(s) containing data relating to a layered mechanism of the respective layered media stream necessary for a party to participate in the layered media stream.

8. A method according to any preceding claim, in which the data contained in a media module includes data necessary for a user to receive or transmit data or both receive and transmit for inclusion in the media session.

9. A method according to any preceding claim, in which the media session is announced by transmitting all of the constituent modules of the session description.

10. A method according to any of claims 1 to 8, in which the media session is announced by transmitting only some of the constituent modules of the session description, with the remaining modules of the session description being subsequently accessible by a user using one or more links provided in the modules transmitted.

11. A method according to claim 10, in which the remaining modules of the session description are held on one or more servers and the one or more links to the remaining modules are in the form of URI pointers.

12. A method according to any preceding claim, in which modules of the session description contain links to modules which are generated subsequent to the announcement.

13. A computer readable storage medium containing data defining at least a part of a description of a media session, the session description comprising:-

a first base module having a first data structure comprising user oriented data relevant to the media session;

at least one media module having a second data structure comprising media oriented data necessary for a user to receive a respective media stream of the media session;

a link between the first base module and the at least one media module;

wherein the link permits a user to access the at least one media module and subsequently receive the media stream.

1/7

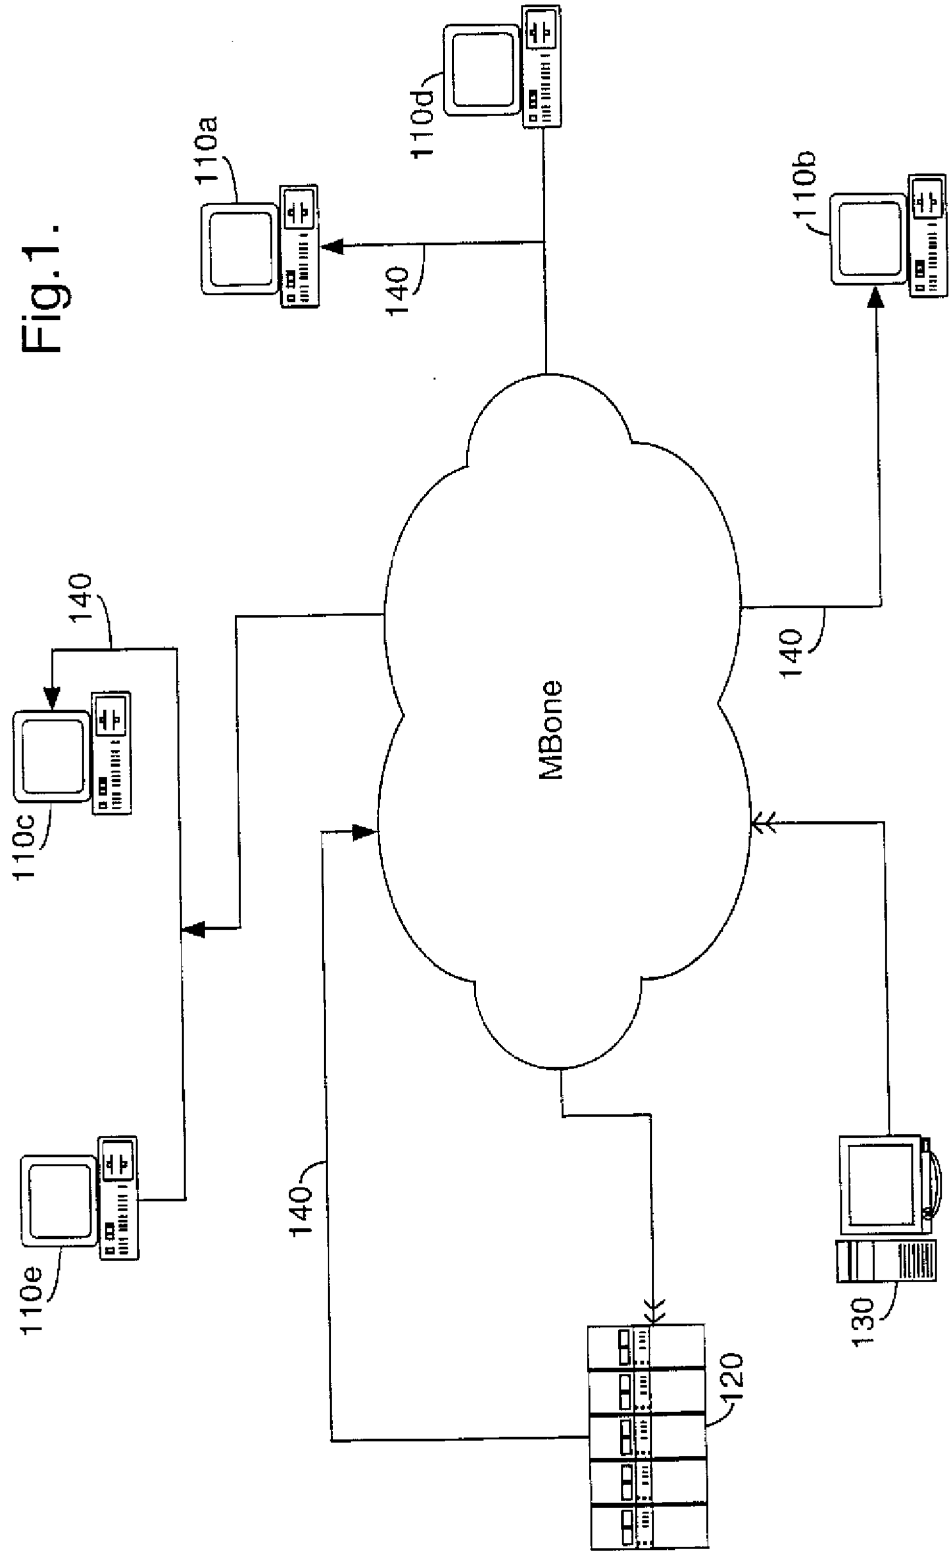

Fig.1.

Fig.2.

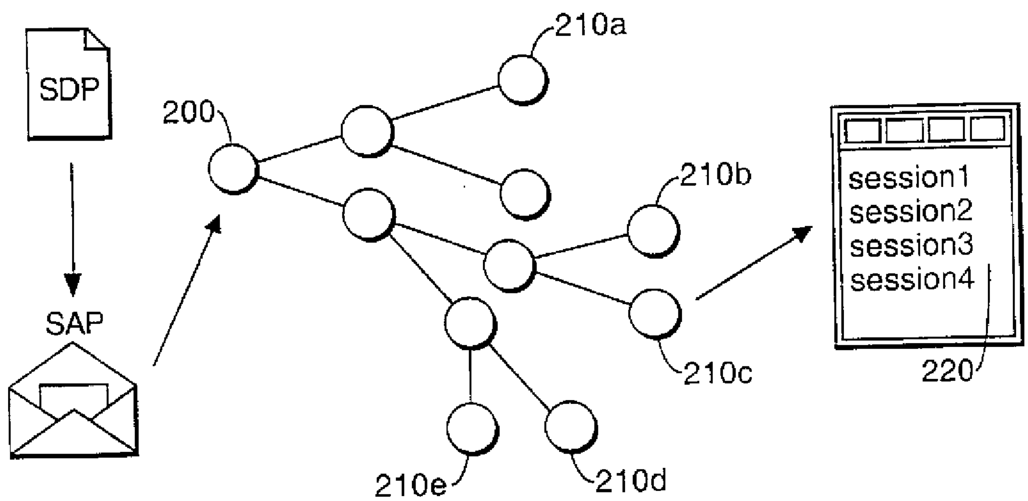

SDP
SAP
200
210a
210b
210c
210d
210e
session1
session2
session3
session4
220

Fig.3.

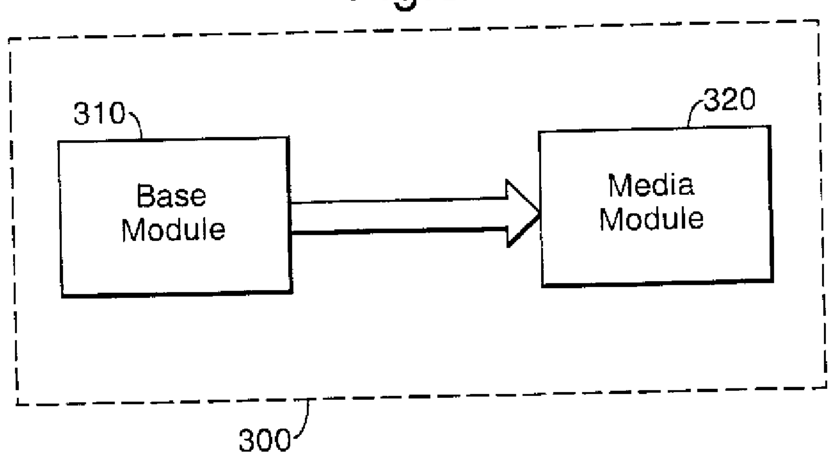

310
Base Module
320
Media Module
300

3/7

Fig.4.

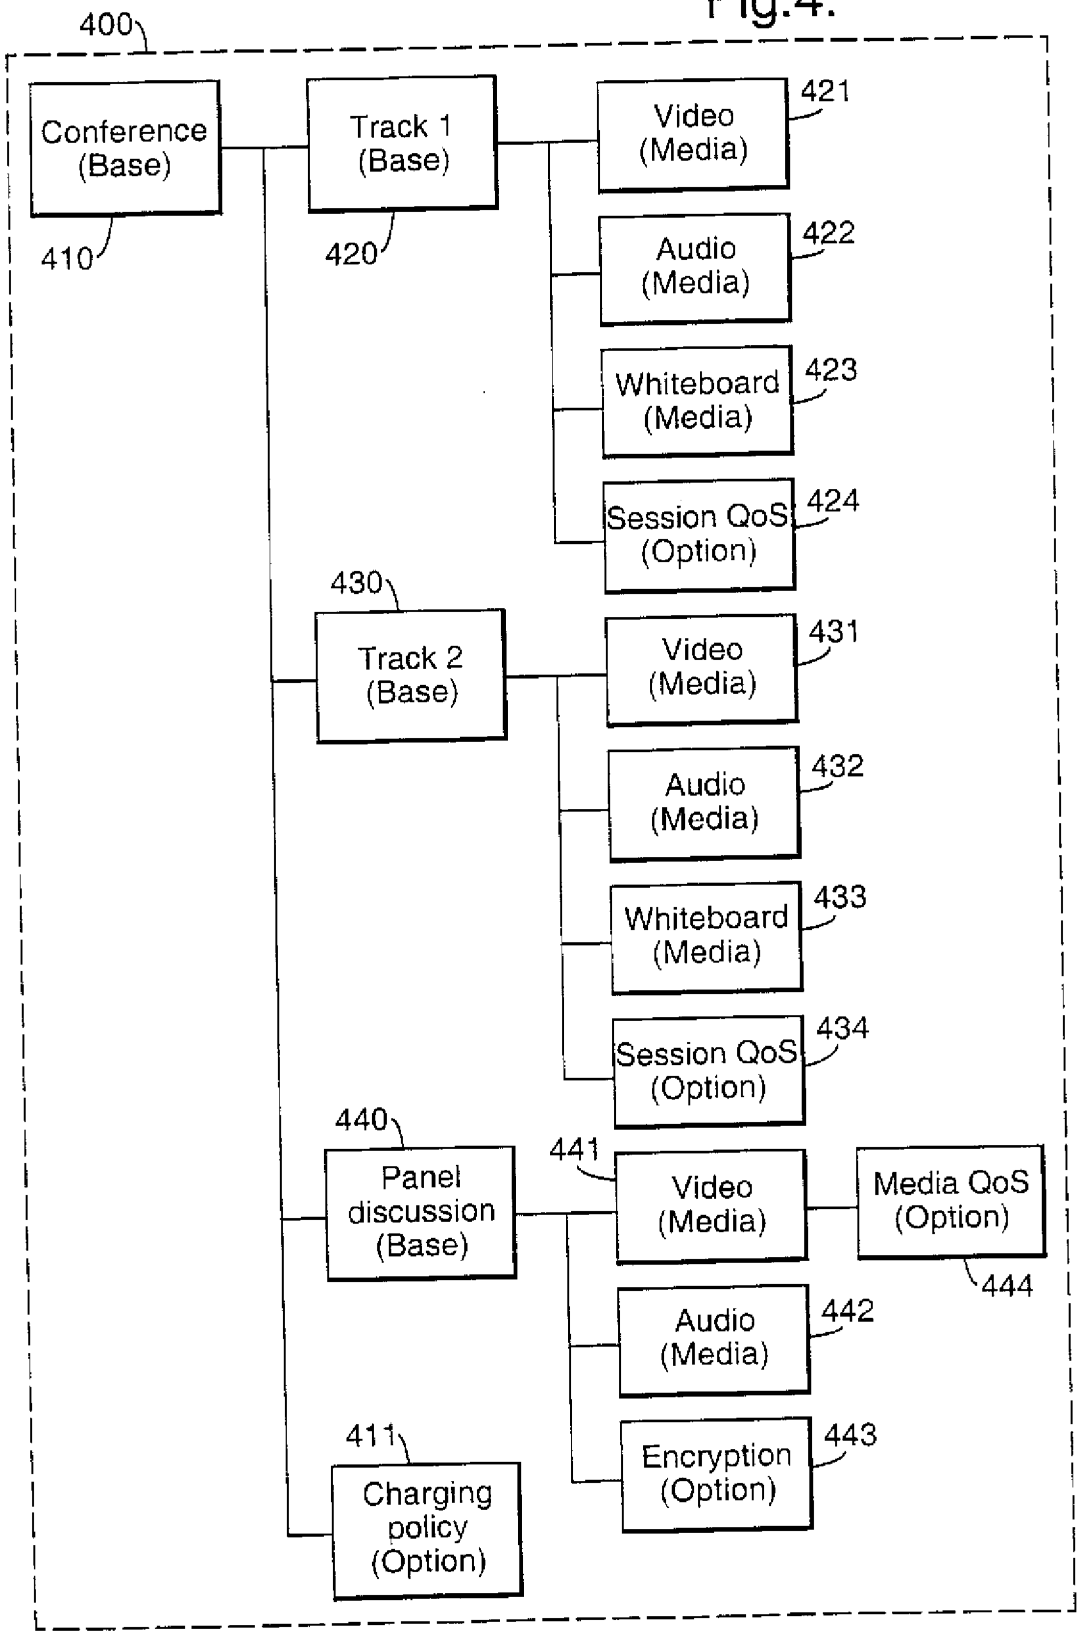
400
Conference (Base)
410
Track 1 (Base)
420
Video (Media)
421
Audio (Media)
422
Whiteboard (Media)
423
Session QoS (Option)
424
430
Track 2 (Base)
Video (Media)
431
Audio (Media)
432
Whiteboard (Media)
433
Session QoS (Option)
434
440
Panel discussion (Base)
441
Video (Media)
Media QoS (Option)
444
Audio (Media)
442
411
Charging policy (Option)
Encryption (Option)
443

Fig.5.

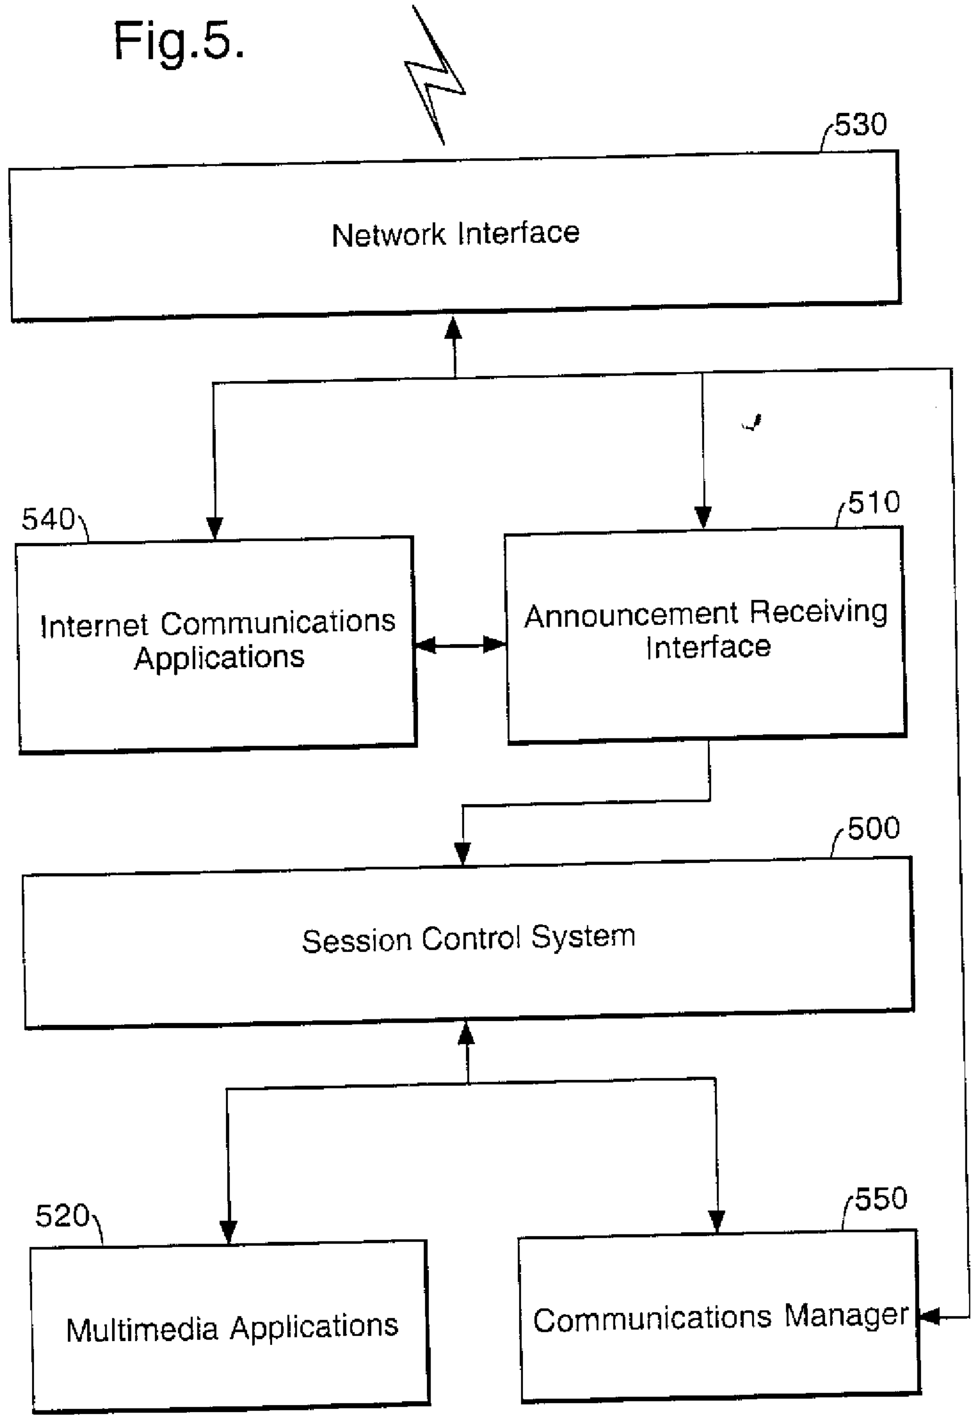
530
Network Interface
540
Internet Communications Applications
510
Announcement Receiving Interface
500
Session Control System
520
Multimedia Applications
550
Communications Manager

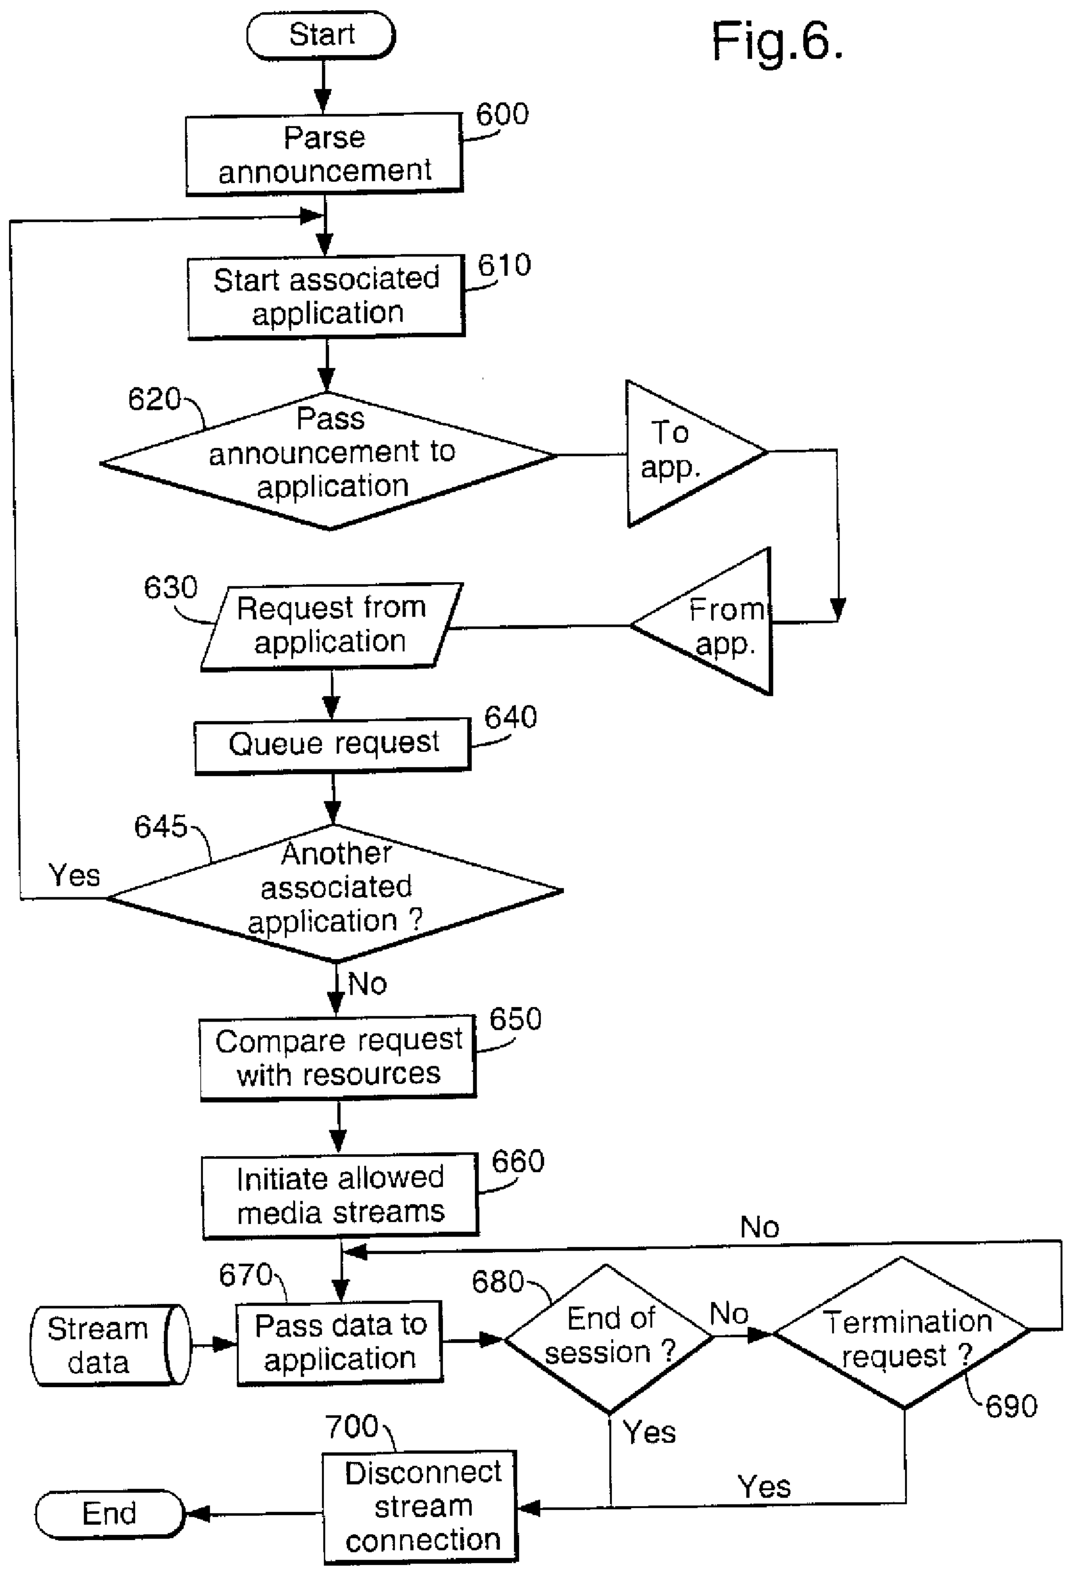
Start
Parse announcement
600
Start associated application
610
620
Pass announcement to application
To app.
630
Request from application
From app.
Queue request
640
645
Another associated application ?
Yes
No
Compare request with resources
650
Initiate allowed media streams
660
No
670
Stream data
Pass data to application
680
End of session ?
No
Termination request ?
690
Yes
Yes
700
Disconnect stream connection
End

Fig.6.

Fig.7.

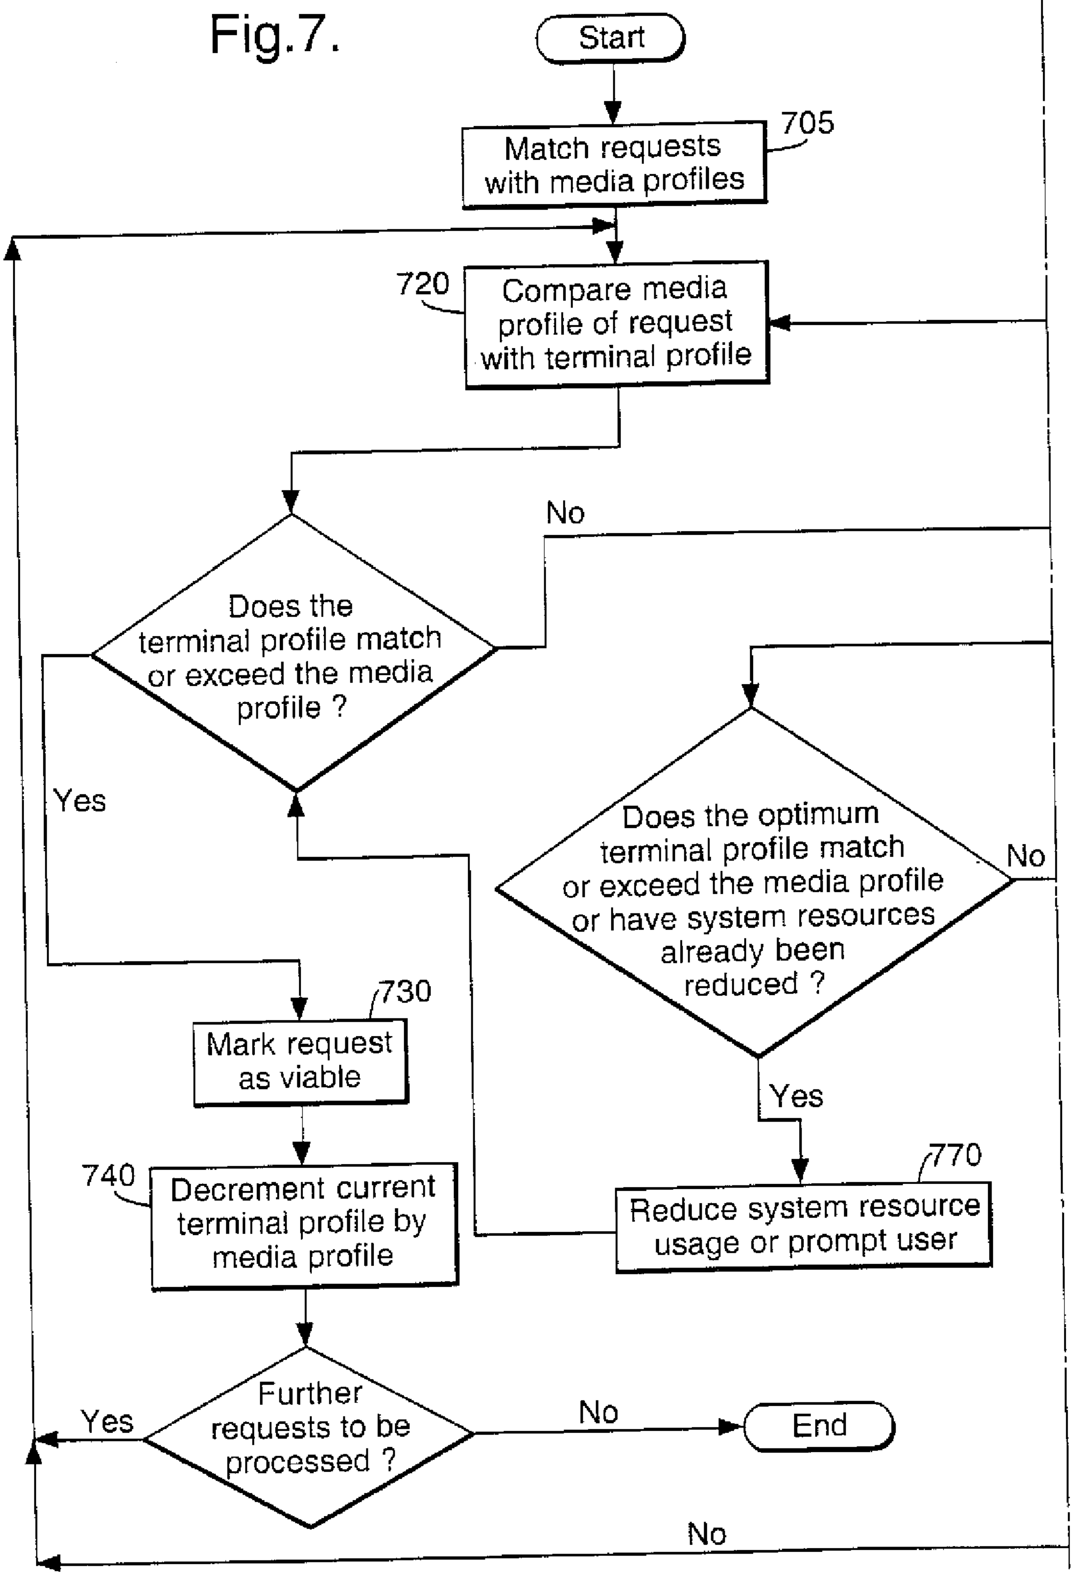
Start
Match requests with media profiles
705
720
Compare media profile of request with terminal profile
Does the terminal profile match or exceed the media profile ?
No
Yes
Does the optimum terminal profile match or exceed the media profile or have system resources already been reduced ?
No
Yes
730
Mark request as viable
740
Decrement current terminal profile by media profile
770
Reduce system resource usage or prompt user
Further requests to be processed ?
Yes
No
End
No

Fig.7 (Cont).

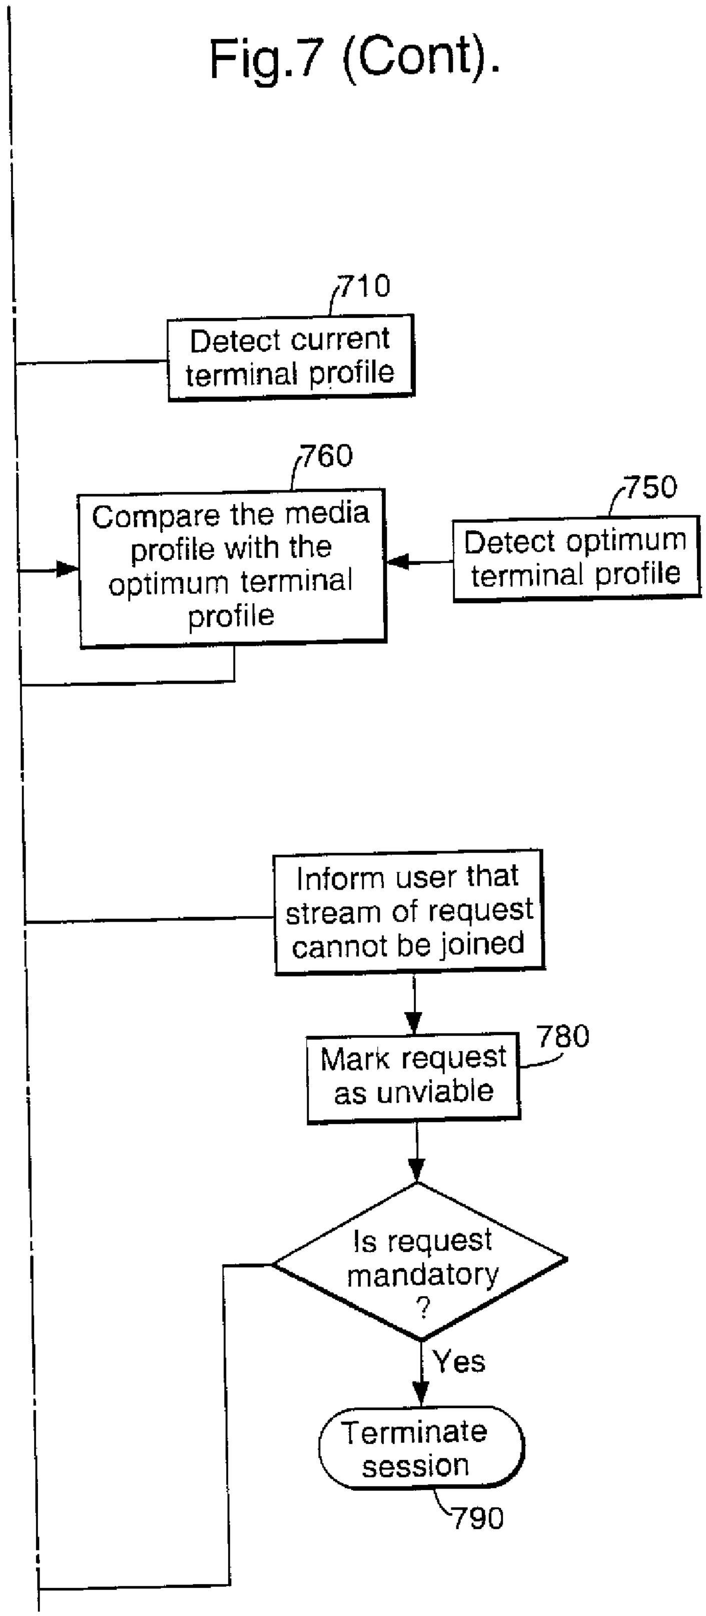
710
Detect current terminal profile
760
Compare the media profile with the optimum terminal profile
750
Detect optimum terminal profile
Inform user that stream of request cannot be joined
Mark request as unviable
780
Is request mandatory ?
Yes
Terminate session
790

INTERNATIONAL SEARCH REPORT

Interr   ıal Application No
PCT/GB 99/03871

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7   H04L29/06   H04L12/18   H04N7/15

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 7   H04L   H04N

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	AVARO O ET AL: "The MPEG-4 systems and description languages: A way ahead in audio visual information representation" SIGNAL PROCESSING. IMAGE COMMUNICATION,NL,ELSEVIER SCIENCE PUBLISHERS, AMSTERDAM, vol. 9, no. 4, 1 May 1997 (1997-05-01), pages 385-431, XP004075337 ISSN: 0923-5965	1,2,8,9, 13
A	abstract page 387, line 17 -page 393, line 2 page 399, line 22-43 page 415, line 6 -page 417, line 7 figure 5 ---	3,7,10, 11

-/--

[X] Further documents are listed in the continuation of box C.

[X] Patent family members are listed in annex.

° Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
20 April 2000	03/05/2000
Name and mailing address of the ISA European Patent Office, P.B. 5818 Patentlaan 2 NL – 2280 HV Rijswijk Tel. (+31–70) 340–2040, Tx. 31 651 epo nl, Fax: (+31–70) 340–3016	Authorized officer Lievens, K

2

Form PCT/ISA/210 (second sheet) (July 1992)

Intern al Application No
PCT/GB 99/03871

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category °	Citation of document, with indication,where appropriate, of the relevant passages	Relevant to claim No.
X	EP 0 798 905 A (DIGITAL VISION LAB CORP) 1 October 1997 (1997-10-01) abstract column 1, line 24 -column 2, line 48 column 4, line 12 -column 6, line 24 column 8, line 49 -column 9, line 5 column 10, line 55 -column 11, line 15 figures 3-5,15 ---	1,8,13
X	THIMM H ET AL: "A MAIL-BASED TELESERVICE ARCHITECTURE FOR ARCHIVING AND RETRIEVING DYNAMICALLY COMPOSABLE MULTIMEDIA DOCUMENTS" MULTIMEDIA TRANSPORT AND TELESERVICES. INTERNATIONAL COST 237 WORKS PROCEEDINGS, VIENNA, NOV. 13 - 15, 1994, 13 November 1994 (1994-11-13), pages 14-34, XP000585292 HUTCHISON D;DANTHINE A; LEOPOLD H; COULSON G (EDS)	1,2,8,13
A	abstract	3,5,6, 10-12
	page 16, line 20 -page 19, line 9 ---	
X	"SDP: Session Description Protocol" RFC2327,April 1998 (1998-04), pages 1-42, XP002101463 http://www.cis.ohio-state.edu/htbin/rfc/rfc2327.html cited in the application	1,9,13
A	page 6, line 20 -page 8, line 6	21,39, 59,60
	page 17, line 15 -page 18, line 38 ---	
A	WO 97 22201 A (XIE DONG ;CAMPBELL ROY H (US); CHEN ZHIGANG (US); TAN SEE MONG (US) 19 June 1997 (1997-06-19) abstract page 9, line 10-15 page 12, line 18 -page 20, line 4 page 36, line 17 -page 37, line 11 -----	1,13

2

INTERNATIONAL SEARCH REPORT

Information on patent family members

Interr   nal Application No
PCT/GB 99/03871

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
EP 0798905 A	01-10-1997	JP 9247142 A	19-09-1997
		CA 2199103 A	05-09-1997
		CN 1170164 A	14-01-1998
		US 5758086 A	26-05-1998
WO 9722201 A	19-06-1997	EP 0867003 A	30-09-1998

PCT WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 7 : H04L 12/46, 12/18	A1	(11) International Publication Number: WO 00/56018 (43) International Publication Date: 21 September 2000 (21.09.00)

(21) International Application Number: PCT/IB00/00150

(22) International Filing Date: 11 February 2000 (11.02.00)

(30) Priority Data:
60/124,066 12 March 1999 (12.03.99) US

(71) Applicant: NORTEL NETWORKS EUROPE S.A. [FR/FR]; 33, quai Paul Doumer, Paris La Defense, F-92415 Courbevoie Cedex (FR).

(72) Inventors: WIGET, Marcel; Villa Azzura, 5, Boulevard du Cap, F-06600 Antibes (FR). PLUIM, Robert; Port Azur II, Apt. I 201, 79 Av. des Freres Roustan, F-06220 Golfe Juan (FR). BRYDEN, Simon; Clos de Ferrogniere, 129 Chemin du Lac, F-06550 La Roquette sur Siagne (FR). MATTSON, Geoffrey; 2, rue Due Reverly, F-06600 Antibes (FR).

(81) Designated States: AE, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CR, CU, CZ, DE, DK, DM, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

Published
With international search report.
Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.

(54) Title: MULTICAST-ENABLED ADDRESS RESOLUTION PROTOCOL (ME-ARP)

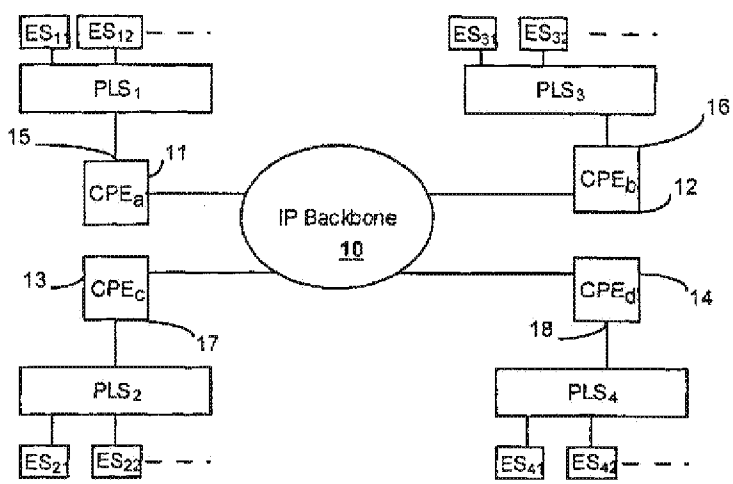

(57) Abstract

A Multicast-Enabled Address Resolution Protocol (ME-ARP) is disclosed. This ME-ARP allows the building of independent IP based Virtual Private LAN segments (VPLS) over a multicast enabled IP backbone using stateless tunnels and optimal VPLS traffic forwarding. Each VPLS has an associated IP subnet which is completely independent from other VPLS or the underlying IP backbone itself. Each Customer Premises Equipment (CPE) device needs only to be configured with a VPLS identifier and its serving IP subnet per VPLS designated interface.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav	TM	Turkmenistan
BF	Burkina Faso	GR	Greece		Republic of Macedonia	TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's	NZ	New Zealand		
CM	Cameroon		Republic of Korea	PL	Poland		
CN	China	KR	Republic of Korea	PT	Portugal		
CU	Cuba	KZ	Kazakstan	RO	Romania		
CZ	Czech Republic	LC	Saint Lucia	RU	Russian Federation		
DE	Germany	LI	Liechtenstein	SD	Sudan		
DK	Denmark	LK	Sri Lanka	SE	Sweden		
EE	Estonia	LR	Liberia	SG	Singapore		

MULTICAST-ENABLED ADDRESS RESOLUTION PROTOCOL (ME-ARP)

FIELD OF THE INVENTION

This invention relates to a scalable and server-less solution to build Virtual Private LAN Segments (VPLS) based on a multicast enabled IP backbone and more particularly to a Multicast-Enabled Address Resolution Protocol (ME-ARP).

BACKGROUND OF THE INVENTION

The popularity of the Internet is driving requirements for secure and segregated IP interconnection of remote sites. One solution is to use the underlying network supporting virtual connections i.e. Frame Relay or ATM. These virtual connections can be separated by provisioning to form a Virtual Private Network which is Layer 3 protocol transparent. However if the underlying network is IP itself, as is the case with the Internet then IP tunnels can be used to interconnect two or more sites. Any other known layer 2 VPN (Virtual Private Network) solution used in the prior art requires a centralized server where all CPE (Customer Premises Equipment) and IP devices have to be statically or dynamically registered, like LANE (Local-Area-Network Emulation), NHRP (Next-Hop-Routing-Protocol) or Classical IP.

A need exists for building IP based virtual private LAN segments (sharing one IP subnet) with complete transparency regarding TCP/IP, site-independent CPE configuration and with dynamic stateless tunnels to optimally forward unicast traffic based on routing and policy per VPLS. VPLS with different Identifiers can use overlapping IP subnets. With the method of the present invention, a centralized server or a list of CPE devices configured for each VPN is not required.

SUMMARY OF THE INVENTION

One aspect of the present invention is to provide a scalable and server-less solution to build Virtual Private LAN Segments (VPLS).

Another aspect of the present invention is to provide a Multicast-Enabled Address Resolution Protocol (ME-ARP). This invention allows the building of independent IP based Virtual Private LAN segments (VPLS) over a multicast enabled IP backbone using stateless tunnels and optimal VPLS traffic forwarding. Each VPLS has an associated IP subnet which is independent from other VPLS or the underlying IP backbone itself. Each Customer Premises Equipment (CPE) device needs only to be configured with a VPLS identifier and its serving IP subnet per VPLS designated interface. In addition, each end station connected to a Physical LAN Segment (PLS) does not need to be modified in order to be a member of the VPLS. No other configuration parameters e.g. list of CPE devices, their logical or physical locations nor their IP addresses are required. The unique invention is ME-ARP (Multicast Enabled Address Resolution Protocol) including the creation of constructed lower layer address based on VPN (Virtual Private Network) Id and tunnel endpoint. Advantages provided by the method of the present invention include:

a) separation of customer IP address space from either the service provider or another customer determined by policy not to be in the same virtual private network (VPN);

b) capability for a remote site to belong to one or more VPN as long as the VPN policy allows. To provide support for IPv4 based applications at this point;

c) transparent or Routed VPN's (by use of external routers) can be constructed independently or combined with this architecture;

d) due to the use of an underlying IP multicast network to forward VPN broadcast traffic in this solution ,there is no need to provide address or broadcast servers; and

e) VPN traffic forwarding is achieved via stateless and optionally secured tunnels which are optimally routed using the underlying IP network backbone routing architecture.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1a is a block diagram illustrating a physical view of a Virtual Private LAN Segment (VPLS) network for use with the present invention;

FIG. 1b is a diagram illustrating a logical view of the network of FIG. 1a or as would be seen from the customer's perspective;

FIG. 2a illustrates a packet format corresponding to an IPsec Authentication Header (AH) encapsulation with authentication;

FIG. 2b illustrates a packet format corresponding to an IPsec Encapsulating Security Payload (ESP) with authentication privacy;

FIG. 3 illustrates a standard ARP packet format on Ethernet;

FIG. 4 is a block diagram of a IP Backbone network for illustrating the ME-ARP request/reply packet flow according to the present invention;

Fig. 5 is a block diagram illustrating the transfer of ME-ARP packet information between a first and second end station according to the method of the present invention; and

Fig. 6 is a table illustrating the content of the ARP tables at various point during the transfer of ME-ARP packet information.

Similar references are used in different figures to denote similar components.

In order to facilitate the description of the invention, the following abbreviations have been used. The terminology used in this document is based on the definitions proposed by the Internet Engineers Task Force (IETF).

CBT Core Based Tree Multicast Routing Protocol

CPE	Customer Premises Equipment
DVMRP	Distance Vector Multicasting Routing Protocol
GRE	Generic Routing Encapsulation
IGMP	Internet Group Management Protocol
LAN	Local Area Network
MOSPF	Multicast extensions for Open Shortest Path First
PA	Provider Address
PIM	Protocol Independent Multicast
PLS	Physical LAN Segment
VPN	Virtual Private Network
VPLS	Virtual Private LAN
UVIP	Unnumbered VPN Internet Protocol

The term "Client Address" (CA) space or network ranges is used to describe the IP address space used by each VPN customer.

The term "Customer Premises Equipment" (CPE) defines an edge device (e.g., router, etc.), fully managed by the provider, connecting a customers PLS to its VPN.

The term "Provider Address" (PA) space or network ranges is used to describe the provider allocated IP addresses in his IP backbone. (e.g., Tunnel endpoints have an address assigned out of the PA range).

The term "Physical LAN Segment" (PLS) is used in this document to describe a broadcast domain, like a shared or switched ethernet segment, connecting hosts, servers and routers at each site. Without the use of a VPN technology, the scope of these PLS is limited per site.

A Virtual Private LAN Segment (VPLS) is the emulation of a LAN segment using Internet facilities. A VPLS can be used to provide what is sometimes known as a transparent LAN service, which can be used to interconnect multiple CPE nodes. It can be seen as a pure layer 2 bridged VPN solution.

The term virtual private networks (VPN) is widely used as a common description for any kind of network built over another network with limited scope.

The term "Unnumbered VPN IP" (UVIP) interface is used in VPLS to describe the tunnel endpoint connecting a PLS on a first site with all other PLS per VPN. In the scope of the customer's PLS, this interface doesn't need to have an IP address assigned to forward traffic (VPLS is a layer 2 VPN solution). The tunnel endpoint itself must have an IP address assigned, out of the providers address space.

DETAILED DESCRIPTION OF THE INVENTION

In order to take advantage of all the features of the present invention, it is assumed that the providers of IP backbone services are IP multicast capable. Similarly, it is assumed that CPE devices are able to join a multicast group using IGMP. It is not a requirement that all routers in the backbone have multicast capabilities. It is possible to interconnect the CPE devices via a partially meshed or "star-like" multicast backbone, built using a mix of multicast routing protocols and tunnels to interconnect multicast islands. IP multicast is used to forward broadcast and multicast traffic and for IP address resolution, but not for forwarding of unicast traffic.

Referring now to Fig. 1a, we have shown the physical view or service provider's view of a Virtual Private LAN Segment (VPLS). The IP backbone 10 and CPE devices 11, 12, 13 and 14 are managed and typically owned by the service provider. CPE devices 11-14 are typically comprised of routers, whereas each PLS is typically comprised of several IP capable devices such as end stations (ES1, ES2, etc.)

Fig. 1b is a diagram illustrating a logical view of the network of Fig. 1a or as would be seen from the customer's perspective. Whereas in Fig. 1a the CPE devices are visible from the provider's perspective, LAN segments are transparent to the customers as illustrated in Fig. 1b. Similarly, CPE devices which are seen by the service provider are invisible to the customer.

Stateless tunnels or links are used in CPE (Customer Premises Equipment) between connected sites. The remote tunnel endpoint address information is directly mapped into the link layer address. ME-ARP is used for IP address resolution inside a VPLS. As a result, VPN connected IP devices will keep all relevant information about the destination tunnel endpoint and VPN membership in their own address resolution (ARP) table.

Special unnumbered IP LAN interfaces will generate the link layer address based on a configured VPN identifier and dynamically learned tunnel endpoints (via ME-ARP).

Again, as illustrated in Fig. 1a and 1b, a VPLS can span two or more sites, with all IP devices sharing the same IP subnet. The IP address and mask are chosen by the customer without any restrictions in relation to the provider or other customers. The CPE devices, managed by the provider, are transparent to the customer. This type of layer 2 VPN solution possesses the following benefits for the customer:

+ Transparency. No IP addresses must be given to the provider;

+ Flat IP subnet. The VPN can be seen as a VPLS, with transparent support for broadcast protocols like DHCP/BOOTP (Dynamic Host Configuration Protocol / BOOTstrap Protocol), Netbios/IP etc; and

+ Broadcast and Multicast support. The customer can extend the VPN with their own routers and run any routing protocol over the VPN without any coordination with the provider.

Each VPLS has a provider wide unique IP multicast address assigned. A UVIP interface of a CPE device, shown at reference numerals 15, 16, 17 and 18, configured for a particular VPLS, will join the VPN's multicast group by using IGMP. All broadcast traffic is then encapsulated and forwarded to the VPN's IP multicast address. There is therefore no need for a central database to keep track of all UVIP interfaces joining a customer's VPN. This is handled by the IP multicast membership.

In order to forward IP unicast traffic, an enhanced version of proxy ARP is used. The differences from the standard proxy ARP are:

a) all ARP requests matching the customers IP subnet are encapsulated and forwarded to all VPN members by sending them to the VPN's IP multicast address. Note: The CPE device cannot determine, if an IP device is connected to the local physical segment or not.

b) a forwarded ARP request, after decapsulation, will replace the source hardware address (MAC, Media-Access-Control or physical Address) not

with the routers own interface MAC address, but by a calculated address containing the tunnel source IP address, an interface unique VPN Id (e.g. VPN instance Id) and a CPE Id (to avoid loops in case of CPE redundancy).

The result of this "multicast enhanced ARP" (ME-ARP) process is that the customers IP devices will keep all relevant information about the destination tunnel endpoint and VPN membership in their ARP table. There is no overhead involved, if compared to a real physical IP subnet.

Unique VPN Identifier

Each VPN has a unique identifier assigned. For VPLS built of more than two physically separated sites this is a valid IP multicast address. As each VPN has a unique IP multicast Id assigned, IGMP and any multicast capable routing protocol (DVMRP (Distance Vector Multicast Routing Protocol), MOSPF (Multicast Open Shortest Path First), PIM (Protocol Independent Multicast), are used by a configured IP VPN interface connecting a Physical Segment to join the VPNs multicast group.

Individual CPE devices are configured as follows:

Based on the VPLS membership using IP multicast, there is no need for a central VPN membership database or protocol to distribute this information. It is enough to configure a new VPN member (physical segment) in the connecting CPE device. The following minimal information is configured per UVIP (Unnumbered VPN IP) interface:

a) VPN IP multicast Id;

b) IP Network/Mask. Assigned by the customer from the Client Address (CA) space. This information is used to determine the correct VPN, based on either source or destination IP address. This is important to support multi-netting on the same physical interface with many VPNs;

c) Tunnel IP address. This address from the Provider Address (PA) space is used to forward VPN traffic over the IP backbone to the correct tunnel end-point (bound to a VPN interface). The VPN identifier in each encapsulated packet can be used to identify the correct logical UVIP interface inside the CPE device;

d) MAC calculation algorithm. This optional, but recommended, configuration parameter allows the support of different MAC address calculation to prevent possible duplicates.

Referring now to Figs. 2a and 2b, in the preferred embodiment of the invention, depending on the security requirements, three different encapsulation formats can be used: without security, with authentication only or with encryption. The encapsulated methods are based on IPsec tunnel mode [RFC2401...RFC2406]. The IP2 header contains the IP source and destination address from the providers address space (tunnel endpoint IP addresses or address as destination address). The IP1 header is the original IP packet header.

In Fig. 2a, we have shown an IPsec AH encapsulation (with authentication). Fig. 2b shows an IPsec ESP encapsulation (with auth. privacy).

IP multicast and broadcast packets are encapsulated and tagged with the VPN multicast Id in the SPI field of the IPsec AH/ESP header and forwarded to the VPN IP multicast address (equal to VPN multicast Id). All active members of the VPNs multicast group receive the encapsulated packet and forward it to the appropriate VPN's UVIP interface.

Referring now to Figs. 3, we have shown an ARP Request/Reply packet including Ethernet transmission layer. In Fig. 4, we have shown a block diagram of an IP Backbone network and in Fig. 5, we have shown a block diagram illustrating the transfer of packet information between a first and second end station, respectively.

In operation, with reference to Figs. 3, 4, 5 and 6, end station A wants to send an IP packet to end station B on the same logical subnet but connected to different gateways. It is assumed, that the ARP tables 80 and 81 from both end stations are empty. Therefore end station A sends an ARP request 50 to the ethernet broadcast address 51. CPE A, configured with the proper VPN information, checks the source IP address 52 of the ARP request packet 50 against its UVIP interfaces configured on the physical interface. In case of a match, it encapsulates the whole, unmodified, ARP request 50 into an IPsec packet 55 including the VPN identifier 56(equals assigned IP multicast address) and forwards packet 55 to the VPN's multicast address 57 using the configured local IP tunnel-endpoint 58 as source address. CPE A also adds a local ARP entry for end station A in its ARP table 72 for that UVIP interface. (CPE A will forward the ARP request, even if end station B is connected to the same physical network).

All CPEs joining the VPN will receive this encapsulated ARP request, unpack it, and forward out the local UVIP interface with the following modification to the original ARP request 55:

replace the original HW source address 59 (MAC address from end station A) with a calculated MAC address containing the tunnel end-point IP address from CPE A(= source address from the received IPsec packet) and an optional interface unique VPN Id.

This new HW source address 60 is replaced in the ethernet header as well as in the ARP packet 61.

CPE B might add an entry to its ARP table 83 for caching. End station B receives the ARP request 62 and respond to it with a normal ARP reply containing its physical HW MAC address 64 as source in the ethernet header and in the ARP reply packet 65. An ARP entry for end station A with the source MAC address from the ARP request is added on end station B. The ARP table 81 of end station B now contains an entry for end station A with a constructed MAC address containing the tunnel-endpoint IP address and VPN Id. CPE B, configured to listen for constructed MAC addresses, identifies the ARP reply 63 from end station B by checking the source MAC address 64 as well as the source IP address 66 (part of VPN's IP network), encapsulate and forwards the ARP reply 67 directly to the addressed tunnel endpoint (extract tunnel endpoint IP address from destination MAC address).

CPE A decapsulates the ARP reply packet 67, checks the destination or target IP address 68 and replaces the destination or target MAC address 69 with the address found in its local ARP cache, and sends the constructed ARP reply 70 out to end station A on the local attached physical LAN segment. In addition, the source MAC address 71(in the Ethernet header and ARP packet) is replaced with a constructed MAC address 72 containing an optional interface locally unique VPN Id and the IP address of CPE B (where the ARP reply came from).

If the ARP table 82 from CPE A does not contain an entry for end station A, then CPE A will have to send an ARP request out for end station A with end station B's IP address before forwarding the ARP reply packet out to end station A.

Finally, end station A receives the ARP reply packet 70 and builds an entry in its ARP table 80 with an entry for end station B and the MAC address containing the remote tunnel endpoint IP address and VPN Id.

THE INVENTION CLAIMED IS:

1. A method of sending a unicast IP packet from a first end station to a second end station, said first and second end stations being on the same logical subnet and connected to different CPEs, comprising:

receiving said unicast IP packet at a CPE associated with said second end station; and

said CPE associated with said second end station providing said second end station with address resolution information containing mapping information between IP and lower layer physical addresses of said first and second end stations, said lower layer physical addresses being constructed by said CPE and containing VPN membership and physical remote location information such that the constructed lower layer addresses contain enough information for said CPE to forward the packet to the correct remote physical location.

2. A method of sending a multicast IP packet from a first end station to multiple end stations, said first and multiple end stations being on the same logical subnet and connected to different CPEs, comprising:

receiving said multicast IP packet at each CPE;
encapsulating said IP multicast packet; and
forwarding said encapsulated IP multicast packet to a VPN assigned multicast address wherein said IP multicast packet is received by each CPE which has been configured to said VPN.

3. A method as defined in claim 2, wherein said multicast IP packet comprises an IP broadcast packet.

4. A method as defined in claim 2, wherein each of said CPE is configured to said VPN using an IP multicast protocol.

5. A method as defined in claim 4, wherein said IP multicast protocol comprises one of an IGMP, DVMRP, MOSPF, MBGP and PIM multicast protocols.

6. A method of sending an IP packet from a first end station to a second end station, wherein said first and second end stations are on the same logical subnet but connected to different CPEs, comprising:

a) sending from said first end station an ARP request with an ethernet broadcast address;

b) at a first CPE associated with said first end station, intercepting said ARP request packet and verifying the intercepted IP address against a corresponding unnumbered virtual packet network (UV) IP interface;

c) if a match is verified, encapsulating said ARP request into an IPsec packet with a VPN identifier; and

d) forwarding said IPsec packet to a VPN's multicast address using configured local IP tunnel-endpoint as a source address.

7. A method as defined in claim 6, wherein said first CPE further adds a local ARP entry for said first end station in its ARP table for said UVIP interface.

8. A method as defined in claim 7, wherein said encapsulated ARP request is received at each CPE connected to said VPN.

9. A method as defined in claim 8, wherein said ARP request is unpacked, modified and forwarded out of the local UVIP interface when received at said CPE.

10. A method as defined in claim 9, wherein said ARP request is modified at each CPE by replacing the original HW source address with a calculated MAC address containing the tunnel end-point IP address from said first CPE and an interface unique VPN Id thus providing a new HW source address to replace in the ethernet header as well as in the ARP packet itself.

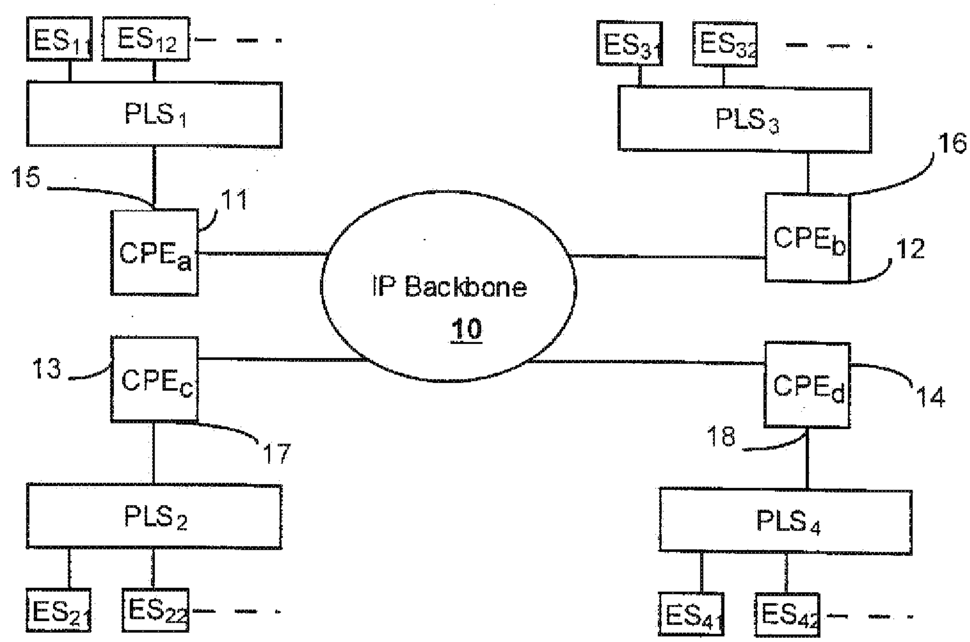
ES11
ES12
ES31
ES32
PLS1
PLS3
15
11
16
CPEa
CPEb
12
IP Backbone
10
13
CPEc
CPEd
14
17
18
PLS2
PLS4
ES21
ES22
ES41
ES42

Fig. 1a

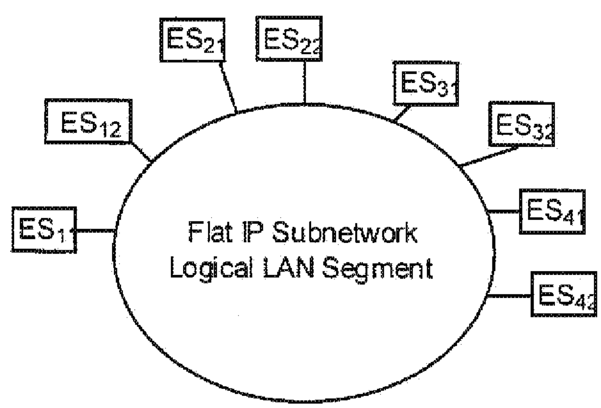
ES21
ES22
ES12
ES31
ES32
ES41
ES11
Flat IP Subnetwork
Logical LAN Segment
ES42

Fig. 1b

IP2 Header	AH Header	IP1 Header	IP1 Payload

Fig. 2a

IP2 Header	ESP Header	IP1 Header	IP1 Payload	ESP Trailer/Auth

Fig. 2b

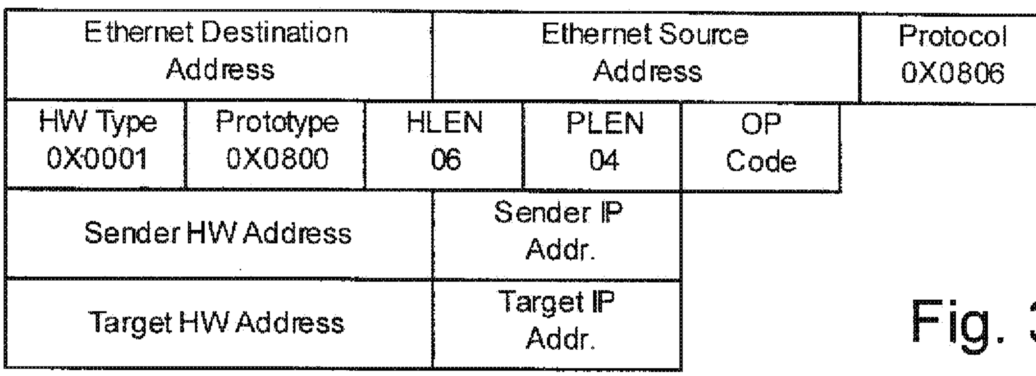

Fig. 3

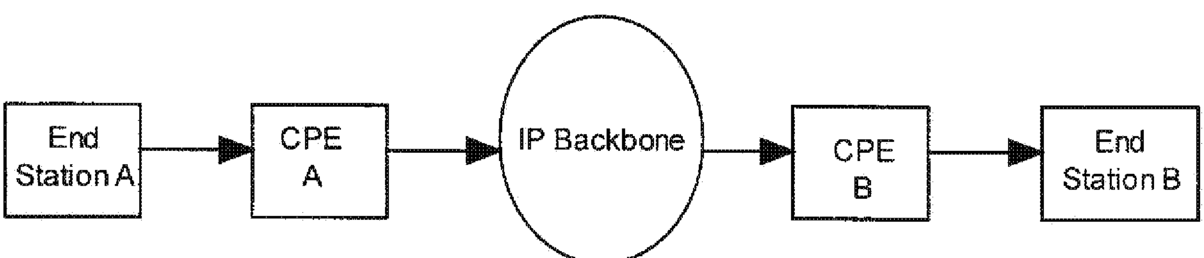

Fig. 4

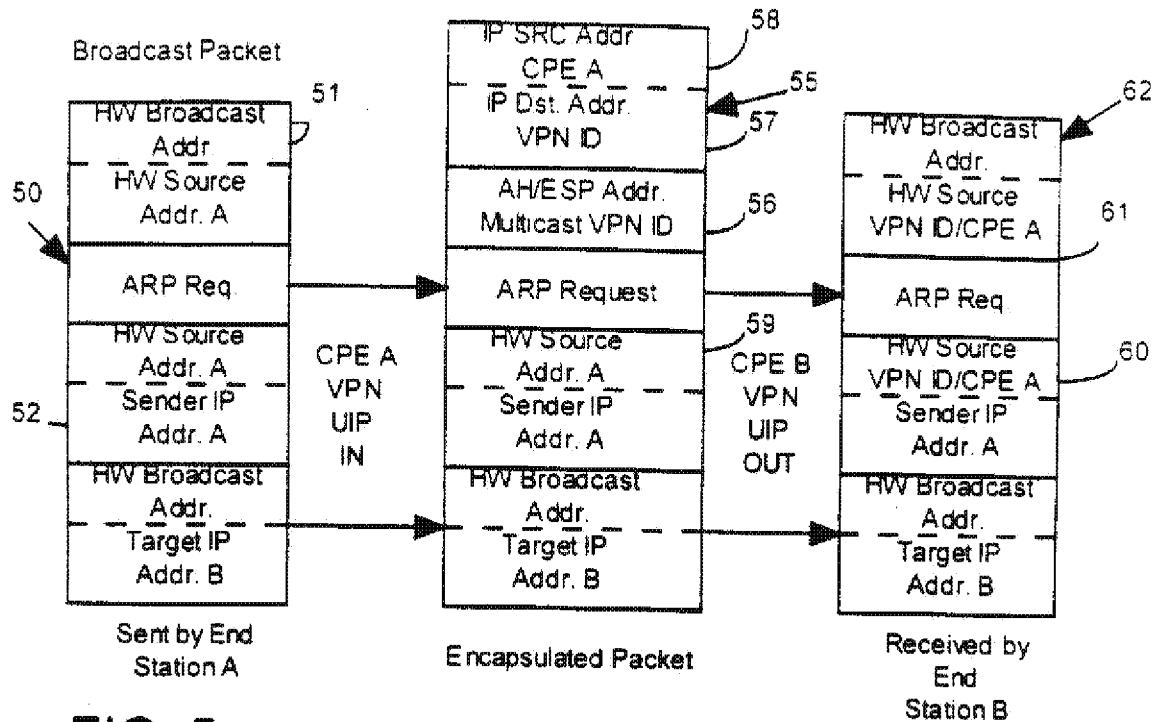
Broadcast Packet
HW Broadcast Addr
HW Source Addr. A
ARP Req.
HW Source Addr. A
Sender IP Addr. A
HW Broadcast Addr.
Target IP Addr. B
Sent by End Station A
50
51
52
CPE A VPN UIP IN
IP SRC Addr CPE A
IP Dst. Addr. VPN ID
AH/ESP Addr. Multicast VPN ID
ARP Request
HW Source Addr. A
Sender IP Addr. A
HW Broadcast Addr.
Target IP Addr. B
Encapsulated Packet
55
56
57
58
59
CPE B VPN UIP OUT
HW Broadcast Addr.
HW Source VPN ID/CPE A
ARP Req
HW Source VPN ID/CPE A
Sender IP Addr. A
HW Broadcast Addr.
Target IP Addr. B
Received by End Station B
60
61
62

FIG. 5

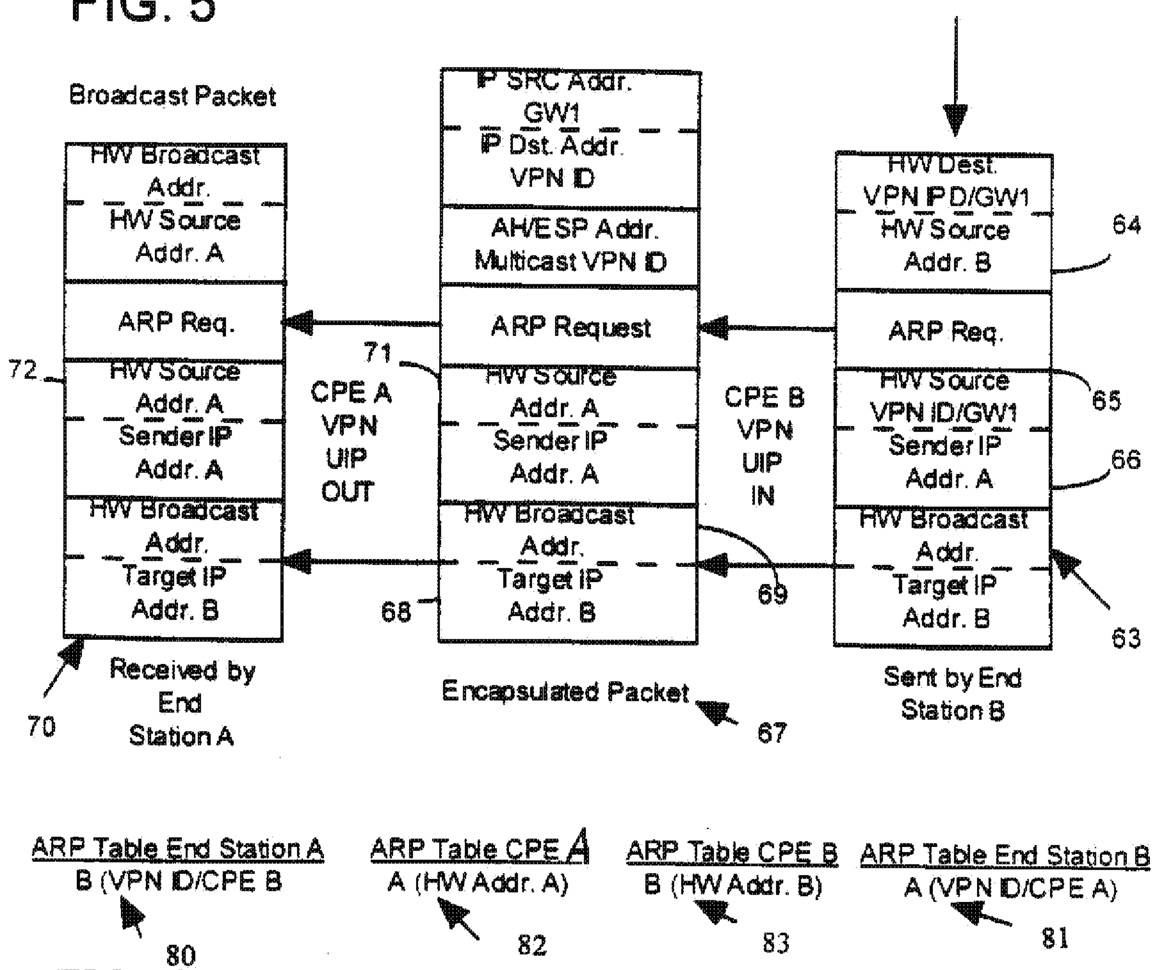
Broadcast Packet
HW Broadcast Addr.
HW Source Addr. A
ARP Req.
HW Source Addr. A
Sender IP Addr. A
HW Broadcast Addr.
Target IP Addr. B
Received by End Station A
70
72
CPE A VPN UIP OUT
71
68
IP SRC Addr. GW1
IP Dst. Addr. VPN ID
AH/ESP Addr. Multicast VPN ID
ARP Request
HW Source Addr. A
Sender IP Addr. A
HW Broadcast Addr.
Target IP Addr. B
Encapsulated Packet
67
69
CPE B VPN UIP IN
HW Dest. VPN IP D/GW1
HW Source Addr. B
ARP Req.
HW Source VPN ID/GW1
Sender IP Addr. A
HW Broadcast Addr.
Target IP Addr. B
Sent by End Station B
63
64
65
66
ARP Table End Station A
B (VPN ID/CPE B
80
ARP Table CPE A
A (HW Addr. A)
82
ARP Table CPE B
B (HW Addr. B)
83
ARP Table End Station B
A (VPN ID/CPE A)
81

FIG. 6

INTERNATIONAL SEARCH REPORT

Inte[rnati]onal Application No
PCT/IB 00/00150

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 H04L12/46 H04L12/18

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 7 H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, WPI Data, PAJ, INSPEC, COMPENDEX, IBM-TDB

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	WO 98 02821 A (3COM CORP) 22 January 1998 (1998-01-22) abstract page 4, line 10 -page 6, line 29 page 13, line 9 -page 15, line 29 ---	1-10
A	EP 0 812 086 A (NIPPON TELEGRAPH & TELEPHONE) 10 December 1997 (1997-12-10) page 2, column 2, line 37 -page 3, column 3, line 35 page 4, column 5, line 2 -page 11, column 50 ---	1-10
A	WO 98 57465 A (VPNET TECHNOLOGIES INC) 17 December 1998 (1998-12-17) page 3, line 13 -page 4, line 8 page 9, line 1 -page 11, line 8 ---	1-10
	-/--	

[X] Further documents are listed in the continuation of box C.

[X] Patent family members are listed in annex.

° Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
28 August 2000	04/09/2000

Name and mailing address of the ISA	Authorized officer
European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Tx. 31 651 epo nl, Fax: (+31-70) 340-3016	Karavassilis, N

1

INTERNATIONAL SEARCH REPORT

Inte. ...onal Application No
PCT/IB 00/00150

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category °	Citation of document, with indication,where appropriate, of the relevant passages	Relevant to claim No.
A	"VIRTUAL PRIVATE NETWORKS ON VENDOR INDEPENDENT NETWORKS" IBM TECHNICAL DISCLOSURE BULLETIN,US,IBM CORP. NEW YORK, vol. 35, no. 4A, 1 September 1992 (1992-09-01), pages 326-329, XP000314784 ISSN: 0018-8689 page 327, line 12 -page 329, column 24 -----	1-10

1

Form PCT/ISA/210 (continuation of second sheet) (July 1992)

INTERNATIONAL SEARCH REPORT

Information on patent family members

Inte[illegible]onal Application No
PCT/IB 00/00150

Patent document cited in search report		Publication date	Patent family member(s)			Publication date
WO 9802821	A	22-01-1998	US	6041166	A	21-03-2000
			AU	3728497	A	09-02-1998
			GB	2330285	A	14-04-1999
EP 0812086	A	10-12-1997	JP	10056473	A	24-02-1998
			US	6075776	A	13-06-2000
WO 9857465	A	17-12-1998	AU	7837998	A	30-12-1998
			EP	09887[illegible]	A	29-03-2000

PCT

WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification : H04L 29/06	A2	(11) International Publication Number: WO 00/51308 (43) International Publication Date: 31 August 2000 (31.08.00)

(21) International Application Number: PCT/SE00/00369

(22) International Filing Date: 25 February 2000 (25.02.00)

(30) Priority Data:
09/257,687 26 February 1999 (26.02.99) US

(71) Applicant: TELEFONAKTIEBOLAGET LM ERICSSON (publ) [SE/SE]; S–126 25 Stockholm (SE).

(72) Inventors: WESTBERG, Lars; Långtora Grän, S–745 96 Enköping (SE). SVANBRO, Krister; Mjölkuddsvägen 133, S–973 43 Luleå (SE). SUNDQVIST, Jim; Regnvägen 80, S–976 32 Luleå (SE).

(74) Agent: MILDH, Christer; Ericsson Radio Systems AB, Ericsson Research/Patent Support Unit, S–164 80 Stockholm (SE).

(81) Designated States: AE, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CR, CU, CZ, DE, DK, DM, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

Published
Without international search report and to be republished upon receipt of that report.

(54) Title: UPDATE OF HEADER COMPRESSION STATE IN PACKET COMMUNICATIONS

(57) Abstract

The soft state of a header compression scheme in a communication system carrying packet traffic including a real time communication signal can be updated (63) during periods of communication signal inactivity (62), during which there is no need to transmit the communication signal. The header compression soft state can also be updated by stealing bits (83, 84) from the communication signal to carry the header update information (73). If the communication signal includes source encoded data, the header compression soft state can be updated selectively (126) based on the bit rate (122, 124) of a codec that produced the source encoded data.

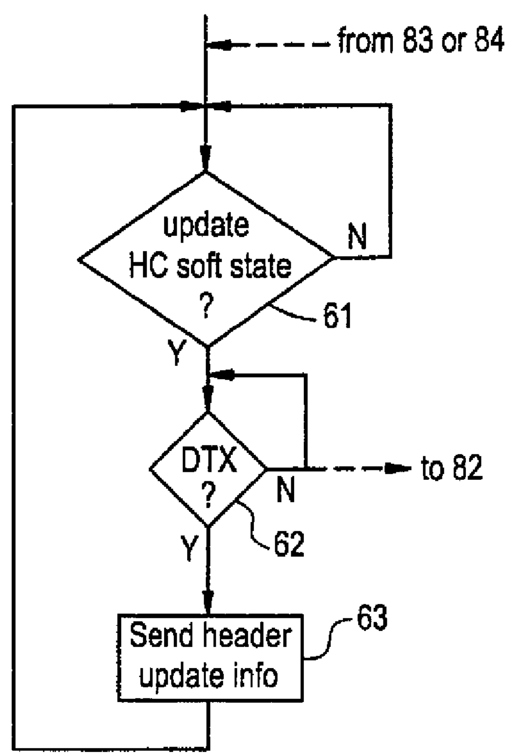

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	**ES**	Spain	**LS**	Lesotho	**SI**	Slovenia
AM	Armenia	**FI**	Finland	**LT**	Lithuania	**SK**	Slovakia
AT	Austria	**FR**	France	**LU**	Luxembourg	**SN**	Senegal
AU	Australia	**GA**	Gabon	**LV**	Latvia	**SZ**	Swaziland
AZ	Azerbaijan	**GB**	United Kingdom	**MC**	Monaco	**TD**	Chad
BA	Bosnia and Herzegovina	**GE**	Georgia	**MD**	Republic of Moldova	**TG**	Togo
BB	Barbados	**GH**	Ghana	**MG**	Madagascar	**TJ**	Tajikistan
BE	Belgium	**GN**	Guinea	**MK**	The former Yugoslav Republic of Macedonia	**TM**	Turkmenistan
BF	Burkina Faso	**GR**	Greece			**TR**	Turkey
BG	Bulgaria	**HU**	Hungary	**ML**	Mali	**TT**	Trinidad and Tobago
BJ	Benin	**IE**	Ireland	**MN**	Mongolia	**UA**	Ukraine
BR	Brazil	**IL**	Israel	**MR**	Mauritania	**UG**	Uganda
BY	Belarus	**IS**	Iceland	**MW**	Malawi	**US**	United States of America
CA	Canada	**IT**	Italy	**MX**	Mexico	**UZ**	Uzbekistan
CF	Central African Republic	**JP**	Japan	**NE**	Niger	**VN**	Viet Nam
CG	Congo	**KE**	Kenya	**NL**	Netherlands	**YU**	Yugoslavia
CH	Switzerland	**KG**	Kyrgyzstan	**NO**	Norway	**ZW**	Zimbabwe
CI	Côte d'Ivoire	**KP**	Democratic People's Republic of Korea	**NZ**	New Zealand		
CM	Cameroon			**PL**	Poland		
CN	China	**KR**	Republic of Korea	**PT**	Portugal		
CU	Cuba	**KZ**	Kazakstan	**RO**	Romania		
CZ	Czech Republic	**LC**	Saint Lucia	**RU**	Russian Federation		
DE	Germany	**LI**	Liechtenstein	**SD**	Sudan		
DK	Denmark	**LK**	Sri Lanka	**SE**	Sweden		
EE	Estonia	**LR**	Liberia	**SG**	Singapore		

UPDATE OF HEADER COMPRESSION STATE IN PACKET COMMUNICATIONS

FIELD OF THE INVENTION

The invention relates generally to packet communications and, more particularly, to header compression in packet communications.

BACKGROUND OF THE INVENTION

The term header compression (HC) refers to the art of minimizing the necessary bandwidth for information carried in packet headers on a per hop basis over point-to-point links. Header compression is usually realized by sending static information only initially. Semi-static information is then transferred by sending only the change from the previous header and completely random information can be sent without compression. Hence, header compression is usually realized with a state machine.

A conventional VoIP-packet (Voice over IP) consists basically of three parts with different quality requirements, as shown in FIGURE 1. The three parts are:

(1) a compressed or not compressed header 11. For example, for real-time speech a conventional IP/UDP/RTP header is often used;

(2) the speech codec bits at part 12, which are most significant for the speech quality. In, for example, the GSM full rate speech codec there are three classes of bits: 1A, 1B and 2, where class 1A and class 2 speech codec bits are respectively most and least important for the speech quality; and

(3) the speech codec bits at part 13 are least important for the speech quality, for example, class 2 bits in GSM.

A conventional header compression scheme for IP/UDP/RTP typically has a soft state characteristic such that the state of the HC may depend on previous headers. An error in a compressed header may result in a loss of the corresponding packet. Because each header usually is represented as a change from the previous header

(delta-coding), an error in a compressed header is a faulty state that will cause successive packets to be lost until the HC soft state is updated. If the payload for the packets with the compressed headers carries a real time service, the loss of several successive packets may be disastrous for the quality of that real time service. For example, the quality of a real time speech service will degrade substantially with successive lost speech frames. If the speech frame error rate has a bursty characteristic, the speech quality will be worse than for the same speech frame error ratio but with a less correlated frame error characteristic.

The effects of bit errors may be different depending on where in the VoIP-packet the bit errors occur:

(1) Bit errors in part 13 of FIGURE 1 (the least important speech codec bits) will result in a slightly degraded quality for the speech carried by that specific packet.

(2) Bit errors in part 12 of FIGURE 1 (the most important speech codec bits) may result in a speech quality degradation so severe that the packet is judged as useless and will not be used in the speech decoder. Hence, that specific packet may be lost due to bit errors in part 12 of the packet.

(3) Bit errors in part 11 of FIGURE 1 (the header, compressed or not) will probably result in the loss of that specific packet since it cannot be transferred to the upper layers of the protocol stack. Further, it can also result in a number of successive lost future packets since the header compression soft state is now corrupt. These are the most severe errors because bit errors in one packet may result in the loss of a number of successive packets.

The conventional header compression algorithms are made for narrow band, wired channels, wherein the error rate of the channel is rather stationary and small. Further, the usage of the channel does not affect other users with similar channels. This is not the case for a wireless channel. The quality of a wireless channel may change rapidly and the usage of the channel affects other users in terms of interference. In a header compression scheme for a wireless channel the probability for errors in the compressed headers will be large and the effect of these compressed header errors has to be reduced.

There are two general approaches to avoid this problem, either minimize the time it takes to update the HC soft state, or minimize the probability for bit errors in compressed headers.

One known way of updating the HC soft state is to send full headers regularly and frequently. For example, a full header can be sent in every fifth speech packet while sending compressed headers in the other packets. If a channel with a fixed bit rate is to be used, the bit rate of this channel is typically chosen with respect to the largest packet size since delay variations are not desirable. Hence, the bit rate of the channel is chosen according to a packet with a full header, resulting in a waste of resources (e.g., radio resources). Further, to achieve robustness in such a header compression scheme, the frequency of full headers must be rather large, which decreases the compression grade and efficiency of the header compression scheme. Hence, regular updates of header compression state with full headers will either result in inefficient header compression or efficient header compression without the necessary robustness against e.g., bit errors.

Another way to update the header compression soft state is for the header compression scheme to demand a soft state update whenever necessary. However, this approach requires a duplex channel with a short round trip time in order to keep the corrupt soft state periods small. Further, such a scheme also requires that the back channel carrying the soft state update request is generally reliable.

It is desirable in view of the foregoing to provide for updating the soft state of a header compression scheme while avoiding the aforementioned disadvantages of prior art approaches.

The present invention provides for updating the soft state of a header compression scheme in a communication system carrying packet traffic including a real time communication signal. The header compression state can be updated during periods when the communication signal is inactive. Also, the invention provides for updating the header compression state by stealing bits from the communication signal to carry the header update information. If the communication signal includes source encoded data, the invention provides for updating the header compression state selectively based on the bit rate of a codec that produced the source encoded data. This operation can permit header compression state updating without stealing any of the source encoded data.

BRIEF DESCRIPTION OF THE DRAWINGS

FIGURE 1 illustrates an exemplary packet format which can be used in conjunction with the present invention.

FIGURE 1A is a shading key for use with FIGURE 1.

FIGURES 2 and 3 illustrate diagrammatically examples of DTX (Discontinuous Transmission) schemes implemented by conventional speech codecs.

FIGURES 4 and 5 illustrate exemplary manners in which the present invention can utilize the conventional DTX operations of FIGURES 2 and 3 to transmit header compression soft state update information.

FIGURE 5A is a shading key for use with FIGURES 2-5.

FIGURE 6 illustrates exemplary operations associated with the header compression update schemes illustrated in FIGURES 4 and 5.

FIGURE 7 illustrates diagrammatically examples of bit stealing operations performed according to the present invention to permit header compression soft state updates.

FIGURE 7A is a shading key for use with FIGURE 7.

FIGURE 8 illustrates exemplary operations associated with the bit stealing scheme of FIGURE 7.

FIGURE 9 illustrates an exemplary packet which can be used in conjunction with the DTX update schemes of FIGURES 4 and 5.

FIGURE 10 illustrates an exemplary packet which can be used in conjunction with the bit stealing scheme of FIGURE 7.

FIGURE 11 illustrates exemplary operations which can be performed in support of HC soft update when receiving packets according to the invention.

FIGURE 12 illustrates pertinent portions of an exemplary communication station according to the invention.

FIGURE 13 illustrates exemplary operations that can be performed in support of HC soft update according to the invention when the packet payload information includes source encoded data.

DETAILED DESCRIPTION

Example embodiments of the invention are cooperable with DTX techniques used in most conventional digital speech services. DTX (Discontinuous Transmission) comprises techniques for detecting non- speech (silent) periods and sending only silence descriptors (SID frames) during these periods in order to produce comfort noise at the receiving end. This comfort noise provides the illusion of continuous transmission of sound. Thus, during non-speech periods, the transmitted packets have a format similar to that shown in Figure 1, except the payload portion (at 12 and 13) includes a SID frame. FIGURES 2 and 3 show conventional DTX schemes, namely the original DTX (FIGURE 2) and the so-called soft DTX (FIGURE 3).

According to an exemplary embodiment of the present invention, header update information can be added to a SID frame of Figure 2 or can replace a SID frame of FIGURE 2. In GSM for example, SID frames (see 21 in FIGURE 2) are transmitted regularly during silent periods (once every 0.48 seconds). The desired update of the header compression state may be accomplished by sending the header update information, for example a full header, together with (see 41) or instead of (see 42) a SID frame, as seen in FIGURES 2 and 4. In another embodiment, the update of header compression state is achieved in conjunction with the conventional soft DTX technique (as described in *"Continuous and Dis-Continuous Power Reduced Transmission of Speech Inactivity for the GSM System"*, Stefan Bruhn et al.,

GlobeCom 98) illustrated in FIGURE 3. The soft DTX technique makes it possible to realize during non-speech periods a low bit rate stream of SID frames 31 which does not introduce much interference to other links. Hence, soft DTX could be used to carry header update information during non-speech periods, as shown in FIGURE 5.

One example of the above-described use of DTX to provide HC soft state updates is shown in FIGURE 6. When an update is desired at 61, it is determined at 62 whether DTX operation is occurring. If so, then the header update information is sent at 63, either in addition to the SID frames (see Figure 5 and 41 of Figure 4) or instead of a SID frame (see 42 in Figure 4).

In conventional video encoding, the transmitting station outputs a sequence of frames that each include, for example, information indicative of a difference between a current captured image and the image captured immediately before the current image. Thus, during periods when the image seen at the transmitting station does not change, the transmitting station sends "static image" frames which indicate that the current image does not differ (or at least does not differ beyond a predetermined limit) from the immediately preceding image. These "static image" frames are thus generally analogous to the aforementioned SID frames, in that they are associated with periods of "static video" wherein no (or no substantial) image change occurs. Accordingly, the techniques described above with respect to FIGURES 2-6 are also applicable to video packet embodiments, the header update information being sent either in addition to the "static image" frames, or instead of a "static image" frame during a period of "static video".

Further exemplary embodiments of the invention replace packet payload bits, e.g., speech frame bits, video frame bits or payload bits representing any desired information, with header compression state update information. If the header compression state is corrupt (e.g., due to bit errors in previous compressed headers) the payload bits (see e.g., 12 and 13 in FIGURE 1) will not be delivered to the application layer until the header compression state is restored. Hence, until the header compression state is restored, the payload bits are useless anyway. Using speech frames as a payload example, by replacing some part of the speech data with header compression update information, immediate future speech frames may be

delivered to the application layer. Parts of a speech frame or the whole speech frame may be replaced with header update information. This replacement of payload bits is also referred to herein as "bit stealing", because payload bits are "stolen" and used instead to carry header update information.

When deciding which speech frame bits to replace with header update information, the characteristics of the speech codec can be taken into consideration. Most conventional speech codecs classify their output bits by relative importance. For example, as mentioned above, the GSM full rate speech codec has three classes of bits with different importance: class 1A, 1B and class 2. Class 1A bits are most important and class 2 bits are least important. Thus, header update information bits would preferably replace class 2 bits where available, because these bits are the least important for the resulting speech quality. FIGURE 7 shows examples of how this can be accomplished.

At 71 in FIGURE 7, all bits except the most important bits are stolen, and all bits are stolen at 72. Considering the updates shown at 73 and 74, fewer bits are stolen for a longer time at 73, while more bits are stolen for a shorter time at 74.

Although the inventive bit stealing techniques of selecting among bits of varying levels of importance are described above with respect to the example of a speech codec that classifies its output bits by relative importance, these bit stealing techniques are applicable to any type of codec that classifies its output bits by relative importance. A video codec is also exemplary of this type of codec.

In embodiments wherein the payload includes source encoded data, the header compression soft state can be updated in conjunction with variations of the bit rate of a codec that produced the source encoded data, and without stealing any of the source encoded data bits. For example, a conventional codec such as a speech or video codec, typically lowers its bit rate for two exemplary reasons: (1) the codec may adapt its bit rate to channel conditions (so-called channel adaptive mode), lowering the bit rate when the channel is congested; and (2) the codec may adapt its bit rate to the behavior of the source (so-called source adaptive mode), lowering its bit rate when the source (for example a speech source or a video source) produces less source stimulus information (i.e., more periods of silence or "static video"). The lowered bit rate in

source adaptive mode is advantageous for sending header update information because less bits are used to represent the source stimulus, leaving more bits to be used for header update information.

FIGURE 13 illustrates exemplary operations that can be performed to implement the above-described use of a lowered codec bit rate to facilitate header compression soft state updates in source encoded data packet embodiments, for example speech or video packet embodiments. When an HC soft state update is desired at 121, it is thereafter determined at 122 whether the codec bit rate is below a threshold level TH. The threshold level TH can be determined empirically to provide desired performance. If the codec bit rate is below TH at 122, then header update information can be sent at 126 in a packet along with the source encoded data.

If at 122 the codec bit rate is not below TH, then it can be determined at 124 whether or not to order the codec to lower its bit rate below TH. If so, then the codec is ordered at 125 to lower its bit rate below TH, and the header update information can be sent at 126 in a packet along with the source encoded data. In embodiments where the codec is not to be ordered to lower its bit rate, operation can flow from 124 back to 122.

After header update information is sent at 126, the codec bit rate can be restored at 127 as needed (i.e., if it was lowered at 125).

The invention also provides for partially updating the header compression state. For example, it may be decided to update only one field (or a few fields) in the header at a given time. As a specific example, if a given speech frame does not have enough bits available for stealing to permit a complete header state update, then perhaps only the RTP sequence number of the RTP portion of an IP/UDP/RTP header would be updated in that speech frame. The use of fewer bits to send partial update information can, in some cases, provide a sufficient HC soft state update but can, in other cases, cause completion of the desired update to take more time (see e.g., 73 in FIGURE 7).

FIGURE 8 illustrates exemplary operations that can be performed to implement a bit stealing scheme according to the invention. If an update is desired at 81, it is determined at 82 whether enough bits are available to be stolen and used to send the complete header update information. If so, then at 83 the bits are stolen and

used to send the complete header update information. If there are not enough bits available at 82, for example, not enough GSM class 2 speech bits, or not enough payload bits in total, then at 84 the available bits are stolen and used to send part of the header update information.

As shown by broken lines in FIGURES 6, 8 and 13, the respective operations shown therein can be variously combined. For example, in speech or video embodiments, if an update is desired in FIGURE 6, but DTX (or "static video") operation is not occurring at 62, then either the bit stealing operations of FIGURE 8 or the codec-related operations of FIGURE 13 can be performed. As another example, if the operations of FIGURE 13 do not result in sending header update information, then either the bit stealing operations of FIGURE 8 or the DTX/"static video" operations of FIGURE 6 can be performed. The decision of whether an update is desired (see 61, 81 and 121) can be made using conventional criteria.

Referring again to the DTX/"static video" update techniques of FIGURES 4 and 5, an example of a packet containing the update information sent during the non-speech/"static video" period is shown in FIGURE 9. The exemplary packet of FIGURE 9 includes a conventional header (compressed or not), a soft state update tag 91, and a header update information portion 93. The soft state update tag 91 makes it possible for a communication station that receives the packet of FIGURE 9 to recognize that the packet includes header update information 93, whereby the receiving communication station will not mistake the FIGURE 9 packet for a conventional speech (or video) packet or a conventional SID (or "static image") frame packet. As shown in broken lines at 94 in Figure 9, the header update information 93 and tag 91 can also be included in a packet with a SID (or "static image")frame, as discussed above with respect to FIGURE 5 and 41 of FIGURE 4.

FIGURE 10 illustrates one example of a packet which can be used to transmit the header update information when using the inventive technique of stealing payload bits and using them to transmit the header update information. The packet of FIGURE 10 includes a conventional header (compressed or not), a soft state update tag 110 and header update information 111. The tag 110 is provided so that a receiving communication station will recognize that the FIGURE 10 packet includes header

update information in addition to (or instead of) payload data. The example of FIGURE 10 indicates in broken lines that a portion 112 of the payload, for example the most significant speech codec bits at 12 of FIGURE 1, can be included in the packet along with the header update information 111.

The packet of FIGURE 10 is also exemplary of a packet that can be used to transmit header update information according to the codec-related technique of FIGURE 13. In this case, the entire payload can be included at 112, because the threshold TH for the lowered codec bit rate can be set as needed to permit the header update information 111 to be added (inserted) without stealing any payload (i.e., source encoded data) bits.

FIGURE 11 illustrates exemplary operations which can be performed according to the present invention in support of HC soft state update when packets are received. After a packet is received at 101, it is determined at 103 whether or not the packet includes a soft state update tag (for example at 91 in FIGURE 9 or 110 in FIGURE 10). If not, there is no HC soft state update. If so, then the header update information (see 93 in FIGURE 9 or 111 in FIGURE 10) is retrieved at 104 and used at 105 to perform the HC soft state update.

FIGURE 12 illustrates pertinent portions of exemplary embodiments of a communication station according to the invention, capable of performing the exemplary operations described above with respect to FIGURES 1-11 and 13. The exemplary communication station of FIGURE 12 can be a wireless station, for example, a mobile radio transceiver such as a cellular telephone, or a fixed-site radio transceiver. The communication station of FIGURE 12 can also be a wireline communication station for use with wired channels, for example a video conferencing host.

The communication station of FIGURE 12 includes a communication port 131 for providing substantive information (for example speech or video information) to a packet unit 132, and for receiving substantive information from the packet unit 132. The communication port 131 also provides header information to a header unit 133. The header unit 133 can use conventional techniques to produce headers (compressed or not) from the header information provided by communication port 131. The header

unit 133 provides outgoing headers to the packet unit 132, and also receives incoming headers from the packet unit 132.

The packet unit 132 is operable conventionally to assemble the header bits received from header unit 133 and the substantive information bits (i.e., payload bits) received from communication port 131 to form an outgoing packet, for example as illustrated in FIGURE 1. The packet unit 132 can forward the assembled packet to a radio unit 134 which transmits the packet over a radio link 135. In other embodiments (e.g. a video conferencing host) the packet unit 132 can output packets to a wired communication channel (e.g. a data network such as the Internet) as shown in broken lines. The outgoing packets in FIGURE 12 can be received by a receiving station (not shown) which can, for example, have structure and functionality analogous to the communication station of FIGURE 12.

The packet unit 132 also receives from the radio unit 134 incoming packets received by the radio unit over the radio link 135. The packet unit 132 conventionally disassembles the incoming packets and provides the substantive information from each incoming packet to the communication port 131 for conventional use. The packet unit also provides the headers from the incoming packets to the header unit 133, which decompresses them as necessary using conventional techniques, and then forwards the header information to the communication port 131.

The packet unit 132 can also receive from the communication port 131 a DTX indication (i.e., no speech activity) or a "static video" indication (i.e., no video activity), to which the packet unit 132 can respond by outputting packets including SID/"static image" frames as illustrated generally in FIGURES 2 and 3.

The packet unit can also communicate with a codec (not shown) to receive therefrom bit rate information and to provide thereto orders to lower/restore the bit rate, as described above with respect to FIGURE 13.

The header unit 133 is coupled to exchange header update information with the packet unit 132, and to signal the packet unit 132 when it is desired to send header update information in an outgoing packet. In response to receiving a signal to send header update information in an outgoing packet, the packet unit 132 can perform the operations illustrated in FIGUREs 6, 8 and 13, either individually or in combination

as desired, as discussed above. A packet such as illustrated at FIGURE 9 can be produced if DTX/"static video" operation is occurring, and a packet such as illustrated at FIGURE 10 can be produced if DTX operation is not occurring.

When the communication station of FIGURE 12 receives an incoming packet, it can perform the exemplary operations illustrated in FIGURE 11. When the packet unit 132 detects an update tag such as illustrated at 91 in FIGURE 9 or 110 in FIGURE 10, the packet unit can retrieve the header update information, and provide this header update information to the header unit 133 along with a signal directing the header unit to update the HC soft state. If, for example, the header update information includes a full header, then the header unit can use the full header in conventional fashion to reset (i.e., update) its header compression state machine (not shown).

It will be evident to workers in the art that the invention described above can be implemented by suitable modifications in hardware, software or both in, for example, a packet communication portion of a conventional wireless or wireline communication station.

As seen from the foregoing discussion, the present invention provides the following exemplary advantages over the prior art: a continuous update of the header compression state may be realized within a constant bit rate channel in a resource efficient way; the time during which the header compression scheme is in a corrupt state is reduced in a resource efficient way; and the number of lost packets due to the corrupt header compression state is reduced, whereby the quality of real-time services is improved.

Although exemplary embodiments of the present invention have been described above in detail, this does not limit the scope of the invention, which can be practiced in a variety of embodiments.

WHAT IS CLAIMED IS:

1. A method of transmitting a communication signal from a first communication station to a second communication station, comprising:

during periods of communication signal activity, sending from the first station to the second station communication signal packets which include header information and communication signal information;

the first station detecting an absence of communication signal activity; and

responsive to the first station detecting an absence of communication signal activity, sending from the first station to the second station an update packet including header update information which can be used by the second station to interpret header information in subsequent communication signal packets sent from the first station to the second station.

2. The method of Claim 1, wherein the communication signal includes one of a speech signal and a video signal.

3. The method of Claim 2, wherein said update packet includes comfort noise information for creating at the second station an illusion of continuous transmission of sound.

4. The method of Claim 2, wherein said update packet sending step includes sending the update packet instead of a packet including comfort noise information.

5. The method of Claim 1, wherein said sending steps include sending the packets via a communication link including a wireless communication channel.

6. The method of Claim 1, including the second station using the header update information to update a header compression state maintained in the second station.

7. The method of Claim 1, including, during one of said periods of communication signal activity, the first station replacing at least some of the communication signal information in one of the packets with header update information.

8. The method of Claim 7, wherein said communication signal information includes source encoded data, and further including, during one of said periods of communication signal activity, the first station determining that a bit rate of a codec that produced the source encoded data is below a threshold level, and thereafter the first station inserting header update information in one of the packets without replacing any of the source encoded data.

9. The method of Claim 8, wherein said determining step includes the first station ordering the codec to lower its bit rate below the threshold level.

10. A method of transmitting information from a first communication station to a second communication station, comprising:

the first station assembling packets which include header information and payload information, and sending the assembled packets from the first station to the second station;

said assembling step including the first station assembling an update packet, including replacing at least some payload information with header update information which can be used by the second station to interpret header information in subsequent packets sent from the first station to the second station; and

said sending step including sending said update packet from the first station to the second station.

11. The method of Claim 10, wherein said replacing step includes replacing all of the payload information with header update information.

12. The method of Claim 10, wherein said replacing step includes replacing a first portion of the payload information with header update information, and wherein said sending step includes sending the header update information in said update packet together with a second portion of the payload information.

13. The method of Claim 12, wherein the second portion of the payload information is a relatively more important portion of the payload information than the first portion thereof.

14. The method of Claim 10, wherein the payload information includes one of speech information and video information.

15. The method of Claim 10, wherein said replacing step includes replacing at least some of the payload information with partial header update information which the second station can use to interpret a portion of the header information in subsequent packets.

16. The method of Claim 15, including determining that an amount of payload information that is available to be replaced by header update information is insufficient to accommodate a desired amount of header update information, and said replacing step including replacing at least some of the payload information with the partial header update information in response to the determination of insufficient available payload information.

17. The method of Claim 10, wherein said sending steps include sending the packets via a communication link including a wireless communication channel.

18. A communication apparatus for transmitting a communication signal to a second communication apparatus, comprising:

a packet unit having an input for receiving communication signal information during periods of communication signal activity, and having an output for sending to

the second apparatus communication signal packets including communication signal information and header information;

a header unit coupled to said packet unit for providing thereto said header information and also for providing thereto header update information which can be used by the second apparatus to interpret header information in subsequent communication signal packets sent from said packet unit to the second apparatus; and

said packet unit responsive to an absence of communication signal activity for sending from said output to the second apparatus an update packet including said header update information.

19. The apparatus of Claim 18, wherein the communication signal includes one of a speech signal and a video signal.

20. The apparatus of Claim 19, wherein said update packet includes comfort noise information for creating at the second apparatus an illusion of continuous transmission of sound.

21. The apparatus of Claim 19, wherein said packet unit is operable to send said update packet instead of a packet including comfort noise information.

22. The apparatus of Claim 19, wherein said packet unit is operable to send the packets via a communication link including a wireless communication channel.

23. A communication apparatus for transmitting information to a second communication apparatus, comprising:

a packet unit having an input for receiving payload information, and having an output for sending to the second apparatus packets including payload information and header information;

a header unit coupled to said packet unit for providing thereto said header information and also for providing thereto header update information which can be

used by the second apparatus to interpret header information in subsequent packets sent from said packet unit to the second apparatus; and

said packet unit operable, before sending one of said packets, to replace at least some of the payload information with the header update information.

24. The apparatus of Claim 23, wherein said packet unit is operable to replace all of the payload information in said one packet with header update information.

25. The apparatus of Claim 23, wherein said packet unit is operable to replace a first portion of the payload information of said one packet with header update information, and to send the header update information in said one packet together with a second portion of the payload information.

26. The apparatus of Claim 25, wherein the second portion of the payload information is a relatively more important portion of the payload information than the first portion thereof.

27. The apparatus of Claim 23, wherein the payload information includes one of speech information and video information.

28. The apparatus of Claim 23, wherein said packet unit is operable to replace at least some of the payload information with partial header update information which the second station can use to interpret a portion of the header information in subsequent packets.

29. The apparatus of Claim 23, wherein said packet unit is operable to send the packets via a communication link including a wireless communication channel.

30. A method of transmitting source encoded data from a first communication station to a second communication station, comprising:

the first station assembling source encoded data packets which include header information and source encoded data, and sending the assembled packets from the first station to the second station; and

the first station determining that a bit rate of a codec that produced the source encoded data is below a threshold level, and thereafter assembling an update packet including header information, the source encoded data and header update information which can be used by the second station to interpret header information in subsequent source encoded data packets sent from the first station to the second station.

31. The method of Claim 30, wherein the source encoded data includes one of speech data and video data.

32. The method of Claim 30, wherein said determining step includes the first station ordering the codec to lower its bit rate below the threshold level.

33. A communication apparatus for transmitting source encoded data to a second communication apparatus, comprising:

a packet unit having an input for receiving source encoded data, and having an output for sending to the second apparatus source encoded data packets including source encoded data and header information;

a header unit coupled to said packet unit for providing thereto said header information and also for providing thereto header update information which can be used by the second apparatus to interpret header information in subsequent source encoded data packets sent from said packet unit to the second apparatus; and

said packet unit having an input for receiving information indicating that a bit rate of a codec that produced the source encoded data is below a threshold level, said packet unit responsive to said information for inserting the header update information in one of said source encoded data packets together with the header information and the source encoded data.

34. The apparatus of Claim 33, wherein the source encoded data includes one of speech data and video data.

35. The apparatus of Claim 33, wherein said packet unit includes an output for ordering the codec to lower its bit rate below the threshold level.

FIG.1
PRIOR ART

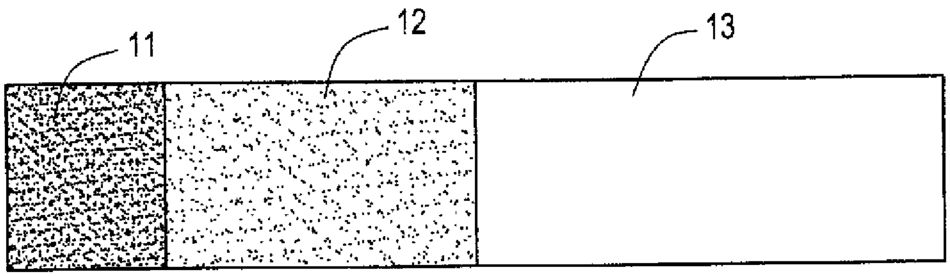

FIG.1A

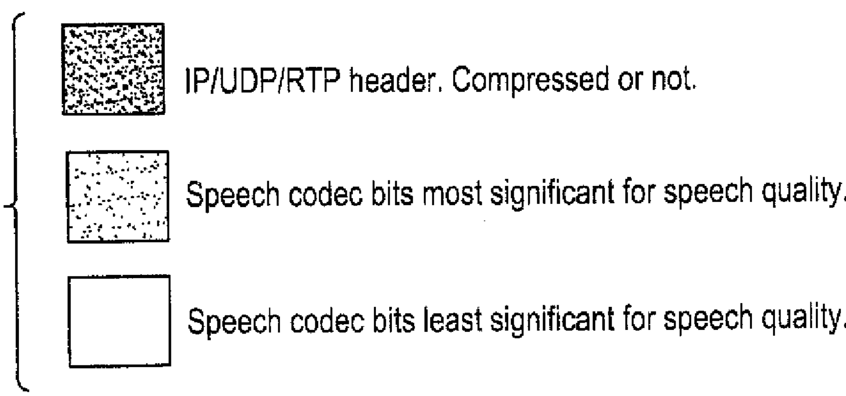

FIG.2
PRIOR ART

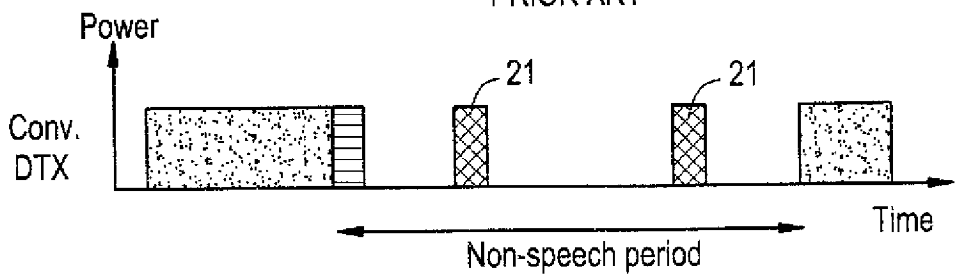

FIG.3
PRIOR ART

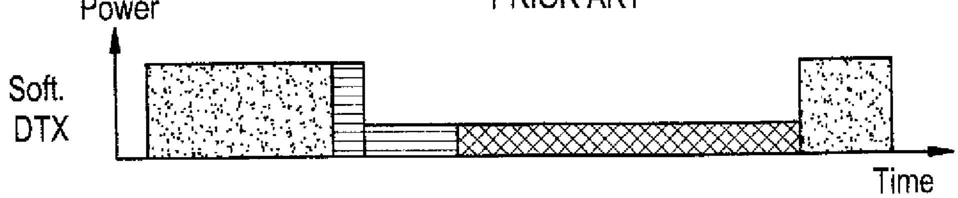

FIG.4

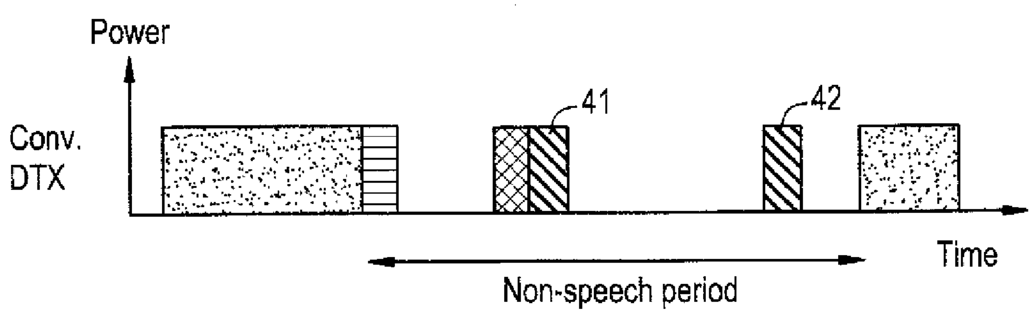

FIG.5

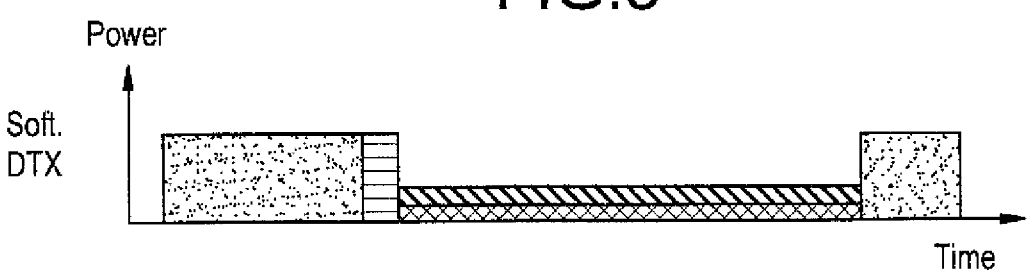

FIG.5A

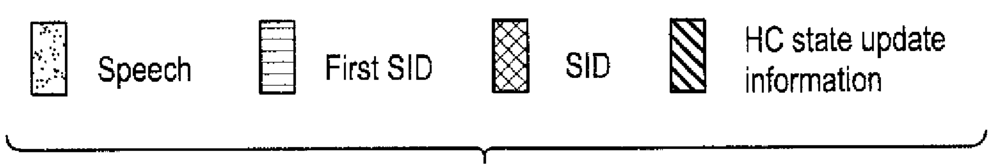

FIG.6

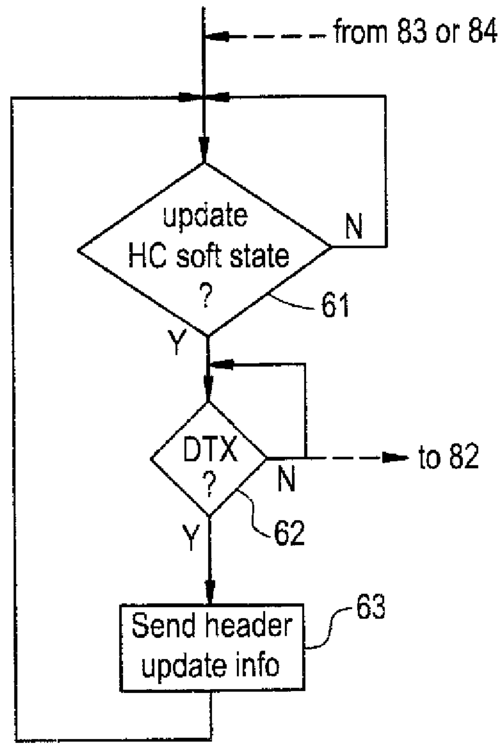

FIG.7

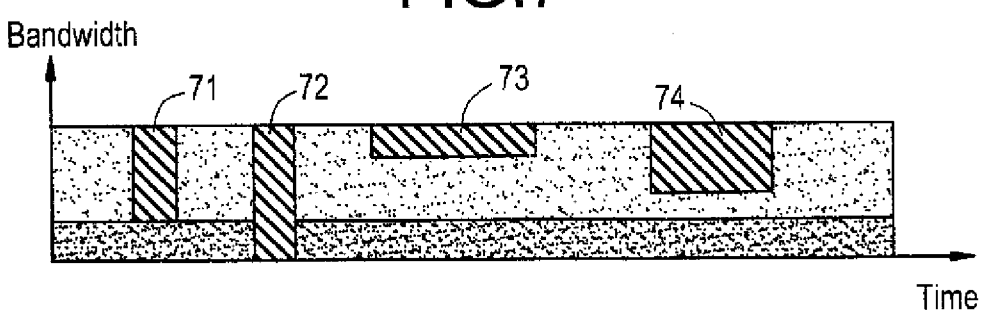

FIG.7A

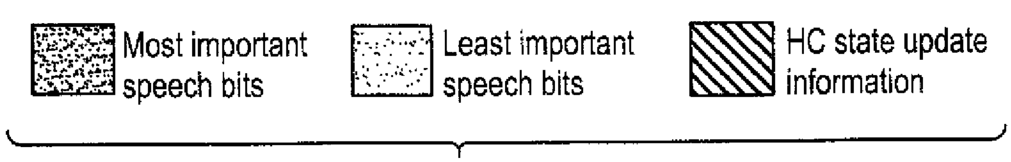

FIG.8

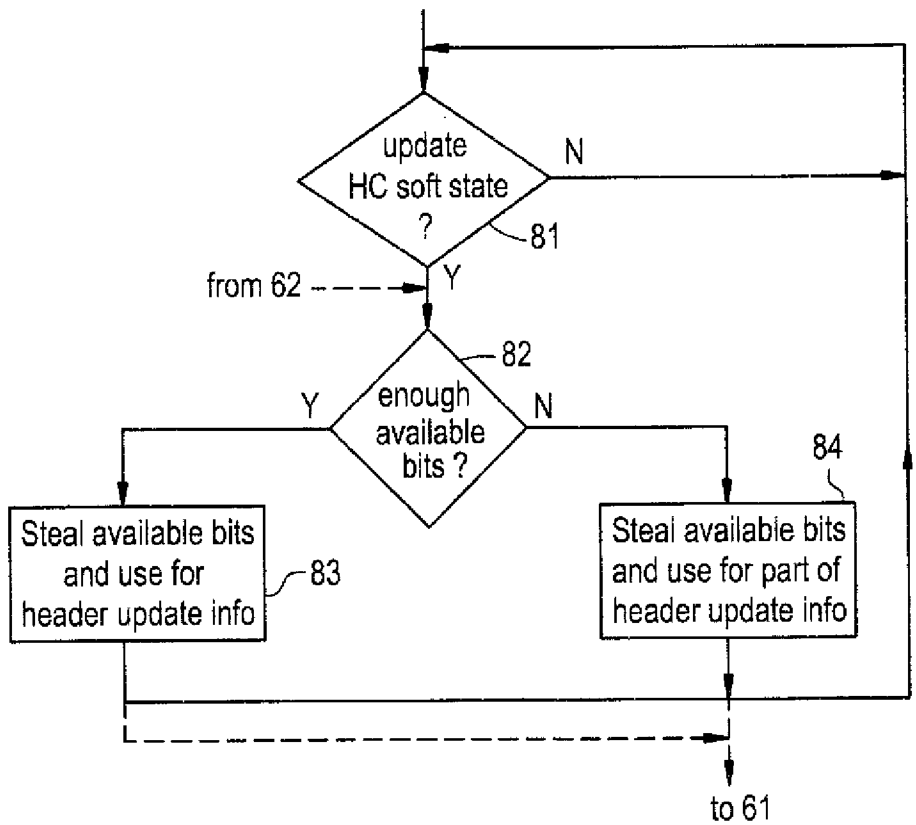
update
HC soft state
?
N
81
from 62
Y
82
enough
available
bits ?
Y
N
84
Steal available bits
and use for
header update info
83
Steal available bits
and use for part of
header update info
to 61

FIG.9

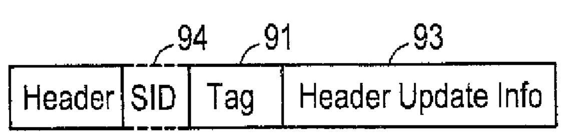
94
91
93
Header
SID
Tag
Header Update Info

FIG.10

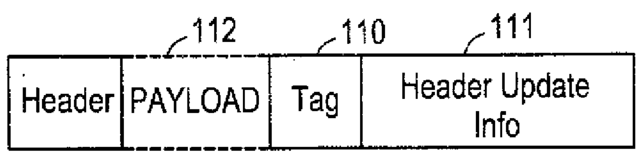
112
110
111
Header
PAYLOAD
Tag
Header Update Info

FIG.11

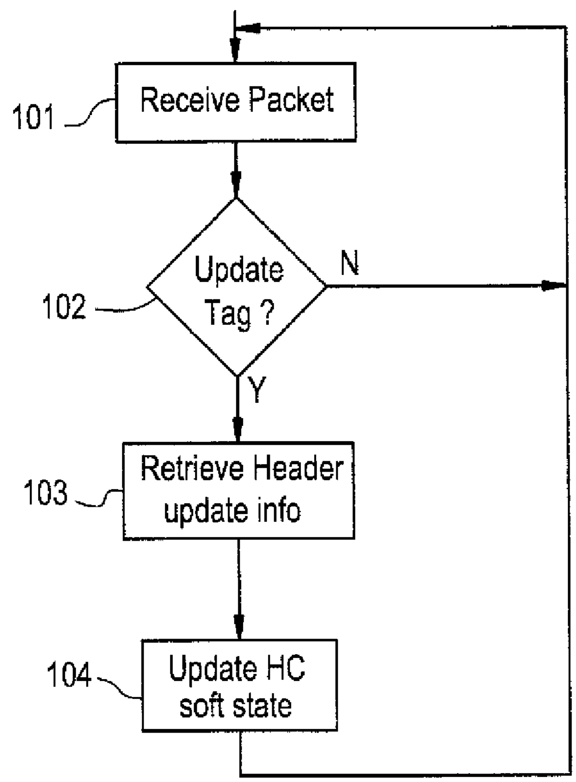
Receive Packet
101
Update Tag ?
102
N
Y
Retrieve Header update info
103
Update HC soft state
104

FIG.12

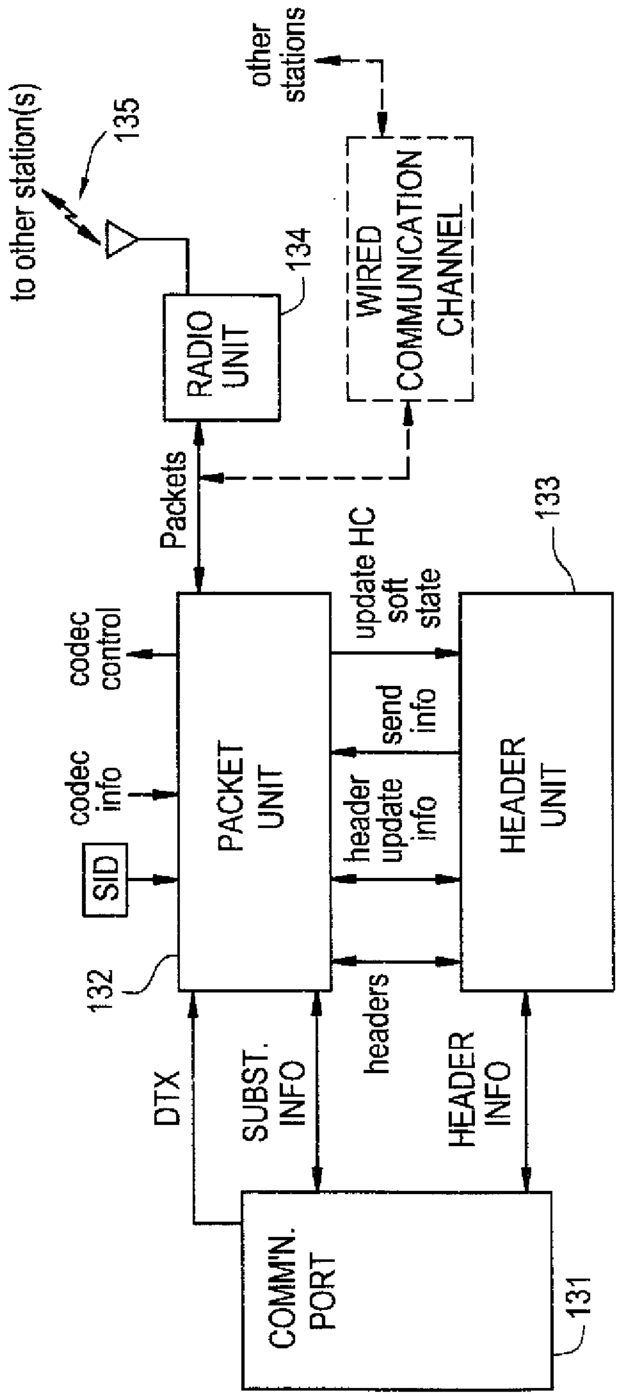
to other station(s)
135
other stations
RADIO UNIT
134
WIRED COMMUNICATION CHANNEL
Packets
codec control
codec info
SID
PACKET UNIT
132
update HC soft state
send info
header update info
headers
133
HEADER UNIT
DTX
SUBST. INFO
HEADER INFO
COMM'N. PORT
131

FIG.13

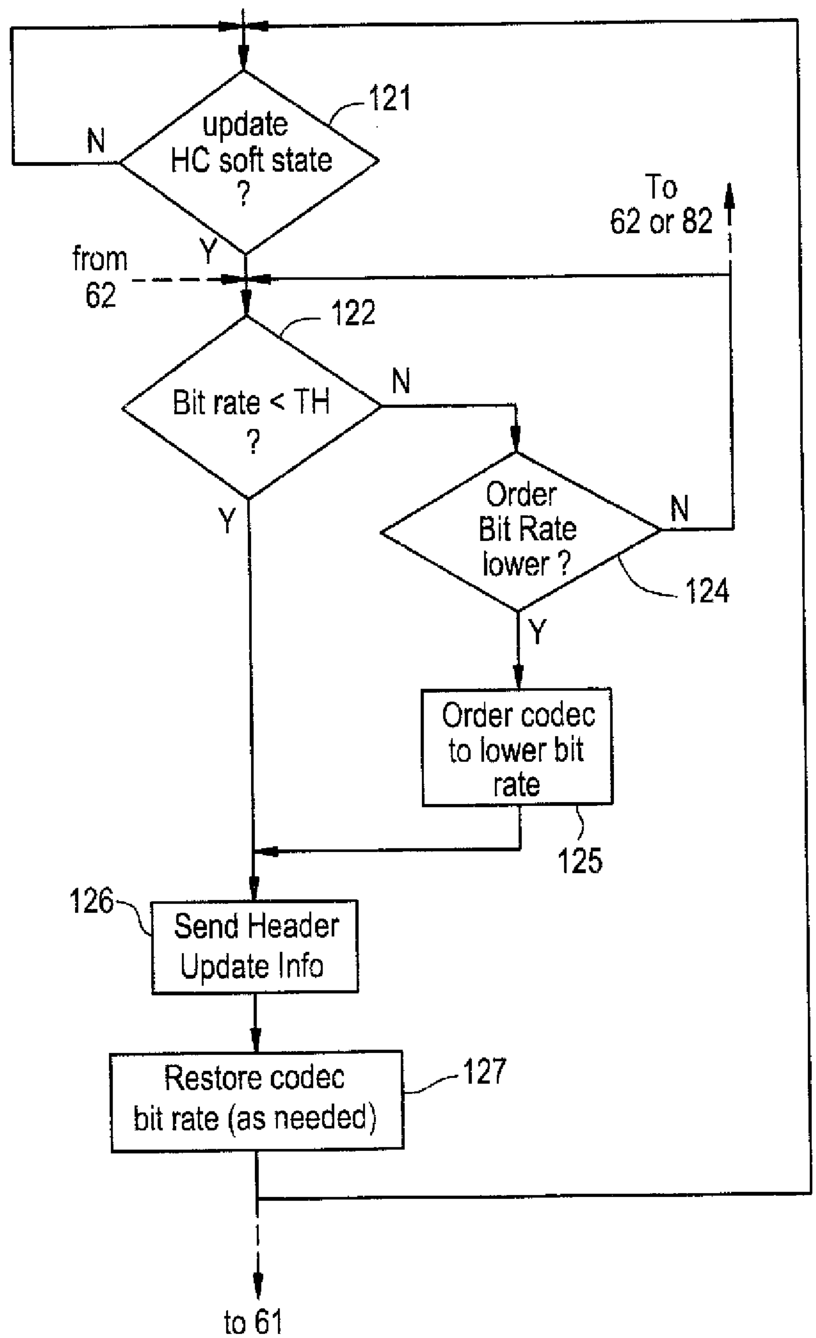

update
HC soft state
?
121
N
Y
from
62
To
62 or 82
122
Bit rate < TH
?
N
Y
Order
Bit Rate
lower ?
N
124
Y
Order codec
to lower bit
rate
125
126
Send Header
Update Info
Restore codec
bit rate (as needed)
127
to 61

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
31 August 2000 (31.08.2000)

PCT

(10) International Publication Number
WO 00/51308 A3

(51) International Patent Classification[7]: H04L 29/06

(21) International Application Number: PCT/SE00/00369

(22) International Filing Date: 25 February 2000 (25.02.2000)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
09/257,687 26 February 1999 (26.02.1999) US

(71) Applicant: TELEFONAKTIEBOLAGET LM ERICSSON (publ) [SE/SE]; S-126 25 Stockholm (SE).

(72) Inventors: WESTBERG, Lars; Långtora Grän, S-745 96 Enköping (SE). SVANBRO, Krister; Mjölkuddsvägen 133, S-973 43 Luleå (SE). SUNDQVIST, Jim; Regnvägen 80, S-976 32 Luleå (SE).

(74) Agent: MILDH, Christer; Ericsson Radio Systems AB, Ericsson Research/Patent Support Unit, S-164 80 Stockholm (SE).

(81) Designated States *(national)*: AE, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CR, CU, CZ, DE, DK, DM, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZW.

(84) Designated States *(regional)*: ARIPO patent (GH, GM, KE, LS, MW, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

Published:
— *With international search report.*

[Continued on next page]

(54) Title: UPDATE OF HEADER COMPRESSION STATE IN PACKET COMMUNICATIONS

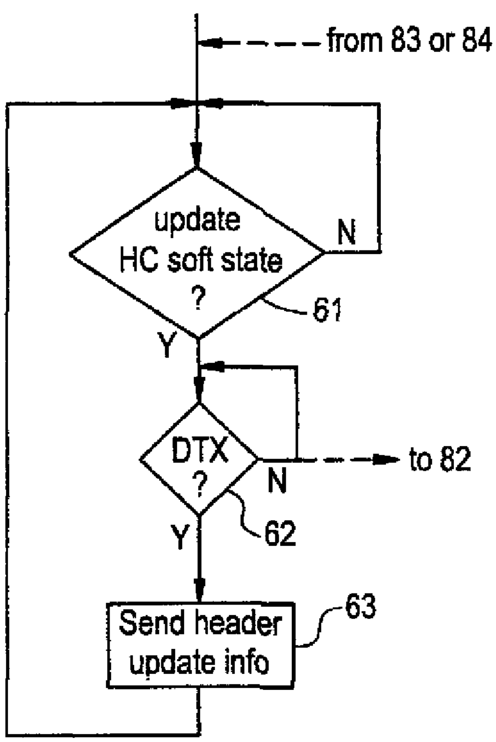

(57) Abstract: The soft state of a header compression scheme in a communication system carrying packet traffic including a real time communication signal can be updated (63) during periods of communication signal inactivity (62), during which there is no need to transmit the communication signal. The header compression soft state can also be updated by stealing bits (83, 84) from the communication signal to carry the header update information (73). If the communication signal includes source encoded data, the header compression soft state can be updated selectively (126) based on the bit rate (122, 124) of a codec that produced the source encoded data.

WO 00/51308 A3

(88) Date of publication of the international search report:
28 December 2000

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

INTERNATIONAL SEARCH REPORT

International Application No
PCT/SE 00/00369

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 H04L29/06

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 7 H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	CASNER S; JACOBSON V: "Compressing IP/UDP/RTP Headers for Low-Speed Serial Links" INTERNET DRAFT, [Online] 27 July 1998 (1998-07-27), XP002125101 Retrieved from the Internet: <URL:ftp://ftp.kyoto.wide.ad.jp/multicast/docs/drafts/draft-ietf-avt-crtp-05.txt> [retrieved on 1999-12-03] page 3, paragraph 3 - paragraph 4 page 16, paragraph 4 -page 19, paragraph 2 --- -/--	1,18

[X] Further documents are listed in the continuation of box C.

[X] Patent family members are listed in annex.

° Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
29 June 2000	19.09.00

Name and mailing address of the ISA	Authorized officer
European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Tx. 31 651 epo nl, Fax: (+31-70) 340-3016	Eraso Helguera, J

1

Form PCT/ISA/210 (second sheet) (July 1992)

International Application No
PCT/SE 00/00369

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	PERKINS S J ET AL: "DEPENDENCY REMOVAL FOR TRANSPORT PROTOCOL HEADER COMPRESSION OVER NOISY CHANNELS" IEEE INTERNATIONAL CONFERENCE ON COMMUNICATIONS (ICC),US,NEW YORK, IEEE, 8 - 12 June 1997, pages 1025-1029, XP000742093 ISBN: 0-7803-3926-6 page 1028, left-hand column, paragraph 1 -right-hand column, paragraph 1 ---	1,18
A	GB 2 294 610 A (FUJITSU LTD) 1 May 1996 (1996-05-01) abstract; figures 5,6 page 2, line 22 - line 33 -----	1,18

1

International application No.
PCT/SE 00/00369

INTERNATIONAL SEARCH REPORT

Box I Observations where certain claims were found unsearchable (Continuation of item 1 of first sheet)

This International Search Report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claims Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:

2. ☐ Claims Nos.:
because they relate to parts of the International Application that do not comply with the prescribed requirements to such an extent that no meaningful International Search can be carried out, specifically:

3. ☐ Claims Nos.:
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box II Observations where unity of invention is lacking (Continuation of item 2 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

1. ☐ As all required additional search fees were timely paid by the applicant, this International Search Report covers all searchable claims.

2. ☐ As all searchable claims could be searched without effort justifying an additional fee, this Authority did not invite payment of any additional fee.

3. ☐ As only some of the required additional search fees were timely paid by the applicant, this International Search Report covers only those claims for which fees were paid, specifically claims Nos.:

4. ☒ No required additional search fees were timely paid by the applicant. Consequently, this International Search Report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

1-9, 18-22

Remark on Protest

☐ The additional search fees were accompanied by the applicant's protest.

☐ No protest accompanied the payment of additional search fees.

FURTHER INFORMATION CONTINUED FROM PCT/ISA/ 210

1. Claims: 1-9,18-22

   Method and communication apparatus for sending from a first station to a second station an update packet including header update information in response to the first station detecting an absence of communication signal activity.

2. Claims: 10-17,23-35

   Method and communication apparatus for sending header update information from a first station to a second station, whereby header update information replaces at least some payload information.

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No
PCT/SE 00/00369

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
GB 2294610 A	01-05-1996	JP 8126047 A	17-05-1996
		GB 2312133 A,B	15-10-1997
		US 5740531 A	14-04-1998

PCT

WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 7 : H04L 12/00 A2

(11) International Publication Number: WO 00/52880

(43) International Publication Date: 8 September 2000 (08.09.00)

(21) International Application Number: PCT/US00/05270

(22) International Filing Date: 29 February 2000 (29.02.00)

(30) Priority Data:

60/122,287	1 March 1999 (01.03.99)	US
60/128,268	8 April 1999 (08.04.99)	US
60/142,082	29 June 1999 (29.06.99)	US
09/505,107	15 February 2000 (15.02.00)	US

(71) Applicant *(for CA only)*: MEDIACELL, INC. [US/US]; 16280 NW Mission Oaks Drive, Beaverton, OR 97006 (US).

(71) Applicant *(for all designated States except CA)*: PACIFIC COMMUNICATIONS LIMITED [CN/CN]; 8/F Bank of America Tower, 12 Harcourt Road, Central, Hong Kong (CN).

(72) Inventors: BISHOP, Donald, M.; 16280 N.W. Mission Oaks Drive, Beaverton, OR 97006 (US). GREEN, George, R., J.; 39078 Guardino Drive #102, Fremont, CA 94538 (US). SHYU, Archie, R.; 12460 Dormouse, San Diego, CA 92129 (US).

(74) Agent: BRUESS, Steven, C.; Merchant & Gould P.C., 3100 Norwest Center, 90 South Seventh Street, Minneapolis, MN 55402-4131 (US).

(81) Designated States: AE, AL, AM, AT, AT (Utility model), AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CR, CU, CZ, CZ (Utility model), DE, DE (Utility model), DK, DK (Utility model), DM, EE, EE (Utility model), ES, FI, FI (Utility model), GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KR (Utility model), KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SK (Utility model), SL, TJ, TM, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

Published

Without international search report and to be republished upon receipt of that report.

(54) Title: METHOD AND APPARATUS FOR COMMUNICATING BETWEEN A CLIENT DEVICE AND A LINEAR BROADBAND NETWORK

(57) Abstract

A method of upstream communication over a linear broadband network includes the steps of generating an upstream baseband signal and modulating it onto an upstream wireless radio frequency carrier to produce a first upstream modulated carrier signal. The modulated carrier signal is transmitted wirelessly, received, and demodulated to reproduce the information integrity of the upstream baseband signal. The signal is then modulated onto an upstream linear broadband radio frequency carrier for transmission on the linear broadband network. Advantageously, noise that accumulates at the subscriber premises is removed from the upstream signal prior to presentation of the signal to the upstream path of the linear broadband network. A system for communicating over a linear broadband network includes network access interface devices coupled to the linear broadband network. A subscriber access interface device accepts upstream communication signals and modulates and transmits the signal to the network access interface device. The network access interface device (6) receives and demodulates the signal and then modulates it for transmission on the linear broadband network.

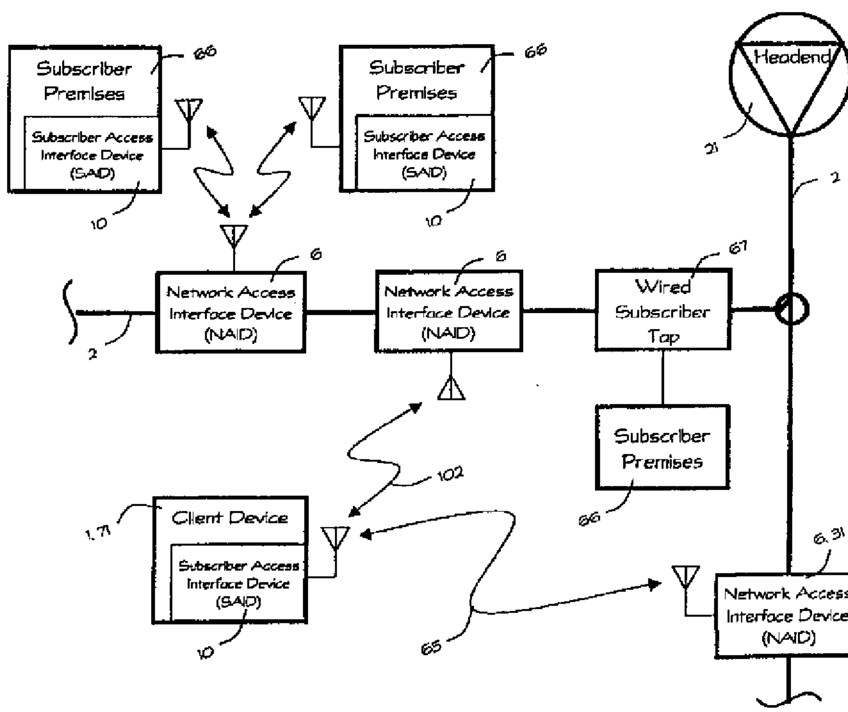

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania
AM	Armenia
AT	Austria
AU	Australia
AZ	Azerbaijan
BA	Bosnia and Herzegovina
BB	Barbados
BE	Belgium
BF	Burkina Faso
BG	Bulgaria
BJ	Benin
BR	Brazil
BY	Belarus
CA	Canada
CF	Central African Republic
CG	Congo
CH	Switzerland
CI	Côte d'Ivoire
CM	Cameroon
CN	China
CU	Cuba
CZ	Czech Republic
DE	Germany
DK	Denmark
EE	Estonia
ES	Spain
FI	Finland
FR	France
GA	Gabon
GB	United Kingdom
GE	Georgia
GH	Ghana
GN	Guinea
GR	Greece
HU	Hungary
IE	Ireland
IL	Israel
IS	Iceland
IT	Italy
JP	Japan
KE	Kenya
KG	Kyrgyzstan
KP	Democratic People's Republic of Korea
KR	Republic of Korea
KZ	Kazakstan
LC	Saint Lucia
LI	Liechtenstein
LK	Sri Lanka
LR	Liberia
LS	Lesotho
LT	Lithuania
LU	Luxembourg
LV	Latvia
MC	Monaco
MD	Republic of Moldova
MG	Madagascar
MK	The former Yugoslav Republic of Macedonia
ML	Mali
MN	Mongolia
MR	Mauritania
MW	Malawi
MX	Mexico
NE	Niger
NL	Netherlands
NO	Norway
NZ	New Zealand
PL	Poland
PT	Portugal
RO	Romania
RU	Russian Federation
SD	Sudan
SE	Sweden
SG	Singapore
SI	Slovenia
SK	Slovakia
SN	Senegal
SZ	Swaziland
TD	Chad
TG	Togo
TJ	Tajikistan
TM	Turkmenistan
TR	Turkey
TT	Trinidad and Tobago
UA	Ukraine
UG	Uganda
US	United States of America
UZ	Uzbekistan
VN	Viet Nam
YU	Yugoslavia
ZW	Zimbabwe

METHOD AND APPARATUS FOR COMMUNICATING BETWEEN A CLIENT DEVICE AND A LINEAR BROADBAND NETWORK

Field of the Invention

The present invention relates generally to data communications and more specifically to transmission and receipt of data via a linear broadband network.

Background of the Invention

There currently exists a complex and robust wired television cable infrastructure that is commonly referred to as the Hybrid Fiber Coax ("HFC") network. The HFC network is an example of a linear broadband network having substantially linear and broadband frequency characteristics. A linear broadband network exhibits linearity in that there are substantially no exponential terms in a gain function of the network over a frequency band of operation. As one of ordinary skill in the art appreciates, an all fiber and an all coaxial cable network is also a linear broadband network. The HFC network merely happens to be the most prevalent linear broadband network in use at the time of the filing of the present patent application. The HFC network has typically been used for delivery of television signals to subscribers. Each subscriber, which represents either an individual or a business, is connected to the cable TV HFC network through coaxial cables running from a headend in a trunk and branch configuration to individual subscribers. Over time, the cable TV HFC network has been upgraded by replacement of some of the coaxial cable trunk lines with fiber optic cable, which has led to this infrastructure being referred to as the HFC network. The connection between the HFC network and the subscriber premises is conventionally made with a coaxial cable, referred to as a subscriber drop, which spans the connection between a tap connected to the HFC network and a client device, which is most cases is a television set, located in the subscriber premises.

Deregulation of the communications industry has made it permissible for the telephony companies to supply television and video services and cable companies to supply telephony and data services. Accordingly, there is an interest among the cable TV service providers to grow their market share by being able to offer all communications services. The cable TV service providers are in a unique position in that they already have a linear and broadband network that reaches many existing subscribers. Their main historical business being television delivery, the cable

companies have focused primarily on the forward or downstream path. In order to be a full service provider, however, the return or upstream path from the subscriber to the headend must be provided. For example, there is a growing demand for communication services that require higher performance from the communication infrastructure, such as higher speed Internet access, interactive television, video conferencing, and telephony. As the demand grows, there will be increasing demands placed on the quality and speed of the downstream and upstream paths. Providing subscriber access to the upstream path presents a challenge to the cable TV service providers. The cable TV service providers have provided a high quality network up to the curb (tap). However, the subscriber drop and client devices have been a source of significant noise resulting from bad connectors, unterminated connections, frayed cables, faulty wiring, breached shielding, noise generated by subscriber appliances, etc. The noise leaks into subscriber wiring and onto the subscriber drop, presenting itself on the upstream path of the HFC network as unwanted signal energy. The very nature of the trunk and branch configuration of the hybrid fiber coax network causes the noise to accumulate on the upstream path as the branches of the network converge. The noise from each subscriber adds together to reduce the overall signal–to–noise ratio of the return signal.

The signal–to–noise ratio of a communications signal is directly related to the effective bandwidth of the channel. Decreasing the signal–to–noise ratio, therefore, increases the bit error rate of a channel. As signal–to–noise ratio decreases, the data transmission rate must slow to a level that provides a sufficiently low bit error rate. The lowering of the data transmission rate is in direct contravention to the objective of the cable TV service providers in supplying high–speed communication services. The signal–to–noise ratio problem is exacerbated when the composite signal reaches an optical laser that is used to power the return transmission fiber. The absolute power level of the signal is limited because the laser has a fixed modulation index. In other words, as the noise level increases, the available signal strength decreases. This limits the cable service provider's option of amplifying the signal to achieve an acceptable signal–to–noise ratio. In data transmission applications, it is possible to employ loss packet retransmission to correct for noise that degrades the integrity of the upstream information. As speeds increase, however, retransmission consumes valuable bandwidth that would otherwise be used for additional upstream information. Consequently, noise limits

the overall capacity of the network, thereby increasing the cost of the service to subscribers. There is a need, therefore, to improve upstream capacity on the network by reducing the injection of upstream noise.

U.S. patent No. 5,867,485 issued to Chambers et al. and assigned to Bellsouth Corporation, proposes a low power microcellular wireless drop for a full duplex interactive network in which a cable connecting a bi–directional fiber network to a subscriber premises is replaced by two wireless transceivers. A Network Interface Unit multiplexes and de–multiplexes signals transmitted and received from a number of subscriber appliances. These signals are transmitted and received by a roof or eaves mounted antenna. The upstream signal is up–converted, amplified, and filtered before being transmitted to a receiver. The system disclosed is a linear processing system, which amplifies the noise presented to the upstream path by the subscriber premises. Disadvantageously, the linear processing propagates any in–band noise and reduces the signal–to–noise ratio. The downstream signal is filtered, amplified, and down–converted before entering the Network Interface Unit and de–multiplexed to the appropriate appliance. The wireless drop succeeds in isolating the subscriber premises from the bi–directional fiber network, but does not remove the noise injected into the upstream signal. There remains a need, therefore, for a method or system to limit the noise ingress into the upstream path.

Summary of the Invention

According to one aspect of an embodiment of the present invention, a method of communicating information from a client device to a linear broadband network having substantially linear and broadband frequency characteristics comprises the steps of generating an upstream baseband signal having a predefined format. The method further comprises modulating the upstream baseband signal onto at least one upstream wireless radio frequency carrier to generate at least one first upstream modulated carrier signal and transmitting the at least one first upstream modulated carrier signal wirelessly. The method further includes receiving the at least one first upstream modulated carrier signal at a network access interface device coupled to the linear broadband network and demodulating the at least one first upstream modulated carrier signal to produce an upstream demodulated baseband signal. The method then comprises modulating the upstream demodulated baseband

signal onto at least one upstream linear broadband radio frequency carrier to produce at least one second upstream modulated carrier signal having a signal format compatible with the linear broadband network.

According to another aspect of the present invention, a method of communicating bi-directional information between a client device and a linear broadband network having substantially linear and broadband frequency characteristics comprises the steps of generating an upstream baseband signal, having a predefined format, and modulating the upstream baseband signal onto at least one upstream wireless radio frequency carrier to generate at least one first upstream modulated carrier signal. The method further comprises transmitting said at least one fist upstream modulated carrier signal wirelessly and receiving said at least one first upstream modulated carrier signal at a network access interface device coupled to the linear broadband network. The method further comprises demodulating the at least one first upstream modulated carrier signal to produce an upstream demodulated baseband signal and modulating the upstream demodulated baseband signal onto at least one upstream radio linear broadband frequency carrier to produce at least one second upstream modulated carrier signal having a signal format compatible with the linear broadband network. The method further comprises receiving at least one downstream linear broadband network radio frequency carrier signal comprising a first downstream modulated carrier signal from the linear broadband network and demodulating the at least one first downstream modulated carrier signal to produce at least one first downstream baseband signal having a predefined format. The method further comprises modulating the at least one first downstream baseband signal onto at least one downstream wireless radio frequency carrier to generate at least one second modulated downstream carrier signal and transmitting said at least one second modulated downstream carrier signal wirelessly. The method further comprises receiving the at least one second modulated downstream carrier signal and demodulating the at least one second modulated downstream carrier signal to produce at least one second downstream baseband signal. The method further comprises transmitting the at least one second downstream baseband signal.

According to another embodiment of a system for upstream communication, the system comprises a linear broadband network having substantially linear and broadband frequency characteristics, a first upstream modulator that modulates at

least one upstream baseband signal received from a client device, the at least one upstream baseband signal being modulated onto at least one upstream wireless radio frequency carrier to produce a first upstream modulated carrier signal. The system further comprises an upstream transmitter that wirelessly transmits the at least one first upstream modulated carrier signal, an upstream receiver that receives the at least one first upstream modulated carrier signal, and an upstream demodulator that demodulates the at least one first upstream modulated carrier signal to generate at least one upstream demodulated baseband signal. The system further comprises, a second upstream modulator that modulates the at least one demodulated upstream baseband signal onto at least one upstream linear broadband network radio frequency carrier for transmission onto the linear broadband network.

According to another embodiment of the present invention a system is provided for bi–directional communication that comprises a bi–directional linear broadband network having substantially linear and broadband frequency characteristics, a first upstream modulator that modulates at least one upstream baseband signal received from a client device, the at least one upstream baseband signal being modulated onto at least one upstream wireless radio frequency carrier to produce at least one first upstream modulated carrier signal, and an upstream transmitter that wirelessly transmits the at least one first upstream modulated carrier signal to an upstream receiver. The upstream receiver receives the at least one first upstream modulated carrier signal. The system further comprises an upstream demodulator that demodulates the at least one first upstream modulated carrier signal to generate at least one upstream demodulated baseband signal. The system further comprises a second upstream modulator that modulates the at least one demodulated upstream baseband signal onto at least one upstream linear broadband network radio frequency carrier for transmission onto the linear broadband network. The bi–directional system further comprises a first downstream receiver that receives at least one first downstream modulated carrier signal from the linear broadband network, a first downstream demodulator that demodulates the at least one first downstream modulated carrier signal to produce a first downstream baseband signal, and a downstream modulator that modulates the first downstream baseband signal onto a downstream wireless radio frequency carrier to produce a second downstream modulated carrier signal. The system further comprises a downstream transmitter that transmits the downstream modulated carrier signal, a second downstream

receiver that receives the downstream modulated carrier signal, and a second downstream demodulator that demodulates the downstream modulated carrier signal to produce a second downstream baseband signal for delivery to the client device.

In another embodiment according to the teachings of the present invention, an apparatus for coupling to a linear broadband network is provided that comprises an upstream receiver that wirelessly receives at least one first upstream modulated carrier signal, an upstream demodulator that demodulates the at least one upstream modulated carrier signal to produce at least one demodulated upstream baseband signal, an upstream modulator that modulates the at least one demodulated upstream baseband signal onto at least one upstream linear broadband network radio frequency carrier to produce at least one second upstream modulated carrier signal for transmission on the linear broadband network, and an upstream transmitter that transmits the at least one second upstream modulated carrier signal onto the linear broadband network.

In another embodiment according to the teachings of the present invention, an apparatus is provided for communicating with a linear broadband network having substantially linear and broadband frequency characteristics comprises an upstream receiver for receiving a plurality of upstream baseband signals over a wired connection from a plurality of client devices, a multiplexer for multiplexing the plurality of upstream baseband signals onto a multiplexed upstream baseband signal, a first upstream modulator for modulating the at least one multiplexed upstream baseband signal onto at least one upstream wireless radio frequency carrier, and an upstream transmitter for transmitting the at least one upstream wireless radio frequency wireless carrier.

In another embodiment according to the teachings of the present invention, a system is provided for upstream communication over a linear broadband network having substantially linear and broadband frequency characteristics, comprises a subscriber access interface device that receives an upstream baseband signal from a client device, modulates said upstream baseband signal onto at least one upstream wireless radio frequency carrier to produce at least one first upstream modulated carrier signal, and wirelessly transmits the at least one first upstream wireless modulated carrier signal, and a network access interface device, coupled to said linear broadband network, that receives the at least one first upstream modulated carrier signal, demodulates the at least one first upstream modulated carrier signal to

produce at least one demodulated upstream baseband signal, modulates the at least one demodulated upstream baseband signal onto at least one upstream linear broadband network radio frequency carrier having a format compatible with the linear broadband network to produce at least one second upstream modulated carrier signal, and transmits the at least one second upstream modulated carrier signal onto the linear broadband network.

In another embodiment according to the teachings of the present invention, a system is provided for upstream communication comprises a linear broadband network having substantially linear and broadband frequency characteristics, a subscriber access interface device that receives an upstream baseband signal from a client device, modulates the upstream baseband signal onto at least one upstream wireless radio frequency carrier to produce at least one first upstream modulated carrier signal, and wirelessly transmits the at least one first upstream wireless modulated carrier signal. The system further comprises a network access interface device, coupled to the linear broadband network, that receives the at least one first upstream modulated carrier signal, demodulates said at least one first upstream modulated carrier signal to produce at least one demodulated upstream baseband signal, modulates said at least one demodulated upstream baseband signal onto at least one upstream linear broadband network radio frequency carrier having a format compatible with the linear broadband network to produce at least one second upstream modulated carrier signal, and transmits the at least one second upstream modulated carrier signal onto the linear broadband network.

In a system employing a method according to the teachings of the present invention, noise that accumulates at the subscriber premises is removed from the upstream signal prior to presentation of the signal to the upstream path of the linear broadband network. Because the upstream signal generated is a digital signal or an analog signal that has been converted to a digital signal, when the carrier is demodulated, the system is able to reconstruct the signal into an image of the transmitted digital signal without the noise, thereby restoring the information integrity of the original signal. The inherent bandwidth limiting function and hysterisis of a digital signal filters out and removes much of the noise, as the signal is prepared for transmission. A baseband signal directly modulates a wireless carrier using modulation techniques that minimize noise interference of the transmitted signal. The reconstruction process further filters out noise that is injected into the

system because of the transmission process. Advantageously, this elimination of noise is cumulative and provides for a significantly quieter upstream path composite signal.

Use of a digital baseband signal for the entire distance up to wireless transmission greatly reduces the introduction of noise and permits use of error checking and retransmission, forward error correction, error concealment, or a combination thereof in the link between the subscriber premises and the network access interface device. Similar use of error checking and retransmission, forward error correction, and error concealment may also be used for the link between the network access interface device and the headend. The additional error management step advantageously reduces the number of errors that must be addressed allowing faster data transmission rates and more efficient use of bandwidth as compared to conventional systems.

As pointed out above, the linear broadband network communications infrastructure may comprise the hybrid fiber coaxial ("HFC") network conventionally used for delivery of cable TV services. The HFC network currently uses a wired tap technology. The wired tap technology comprises a drop cable extending from a tap at a curb to equipment located at a subscriber premises. The wireless tap, or network access interface device according to the teachings of the present invention, is compatible with and is able to coexist on the same network as the wired tap. Accordingly, the network access interface device of the present invention provides a smooth and gradual upgrade installation path for communication service providers. Services to existing subscribers need not be disturbed when subscribers are added or upgraded. Each network access interface device supports a plurality of subscribers. Accordingly, the total number of two–way transceivers connected to the linear broadband network is reduced. By concentrating the services supported by a single modem, there is a decrease in the number of noise sources connected to the linear broadband network. The signal–to–noise ratio, therefore, is improved by 10log the a ratio of the number of network access interface devices over the number of wired modems that a conventional system would require for the same level of service.

Brief Description of the Drawings

Figure 1 is a representative view of an embodiment of system according to the teachings of the present invention.

Figure 2 is a representative view of a network access interface device and various possible embodiments of subscriber access interface devices and client devices employed in a method according to the teachings of the present invention. The possible embodiments shown are not exhaustive. Other embodiments will occur to one of ordinary skill in the art having benefit of the present disclosure.

Figure 3 is a representative view of an upstream signal flow according to the teachings of the present invention from an initiating client device to a linear broadband network.

Figure 4 is a representative diagram of data packets at various stages of an upstream data transmission process.

Figure 5 is a representative view of a downstream signal flow according to the teachings of the present invention from linear broadband network to a destination client device.

Figure 6 is a representative diagram of data packets at various stages of a downstream data transmission process.

Figure 7 is a representative block diagram of an embodiment of a direct sequence spread spectrum network access interface device for implementation of a method according to the teachings of the present invention.

Figure 8 is a representative block diagram of an embodiment of a direct sequence spread spectrum subscriber access interface device for implementation of a method according to the teachings of the present invention.

Figure 9 is a representative block diagram of an embodiment of a frequency hopping spread spectrum network access interface device for implementation of a method according to the teachings of the present invention.

Figure 10 is a representative block diagram of an embodiment of a frequency hopping spread spectrum subscriber access interface device for implementation of a method according to the teachings of the present invention.

Figure 11 is a flow diagram of the polling process of an embodiment according to the teachings of the present invention.

Detailed Description of Embodiments of the Invention

With specific reference to Figures 1 and 2 of the drawings, there is shown a system according to the teachings of the present invention in which a subscriber access interface device 10 ("SAID") communicates bi-directionally with a network access interface device 6 ("NAID") coupled to a linear broadband network 2. The NAID 6 is an integrated gateway that contains both a Wide Area Network ("WAN") connection and a wireless Local Area Network ("LAN") connection. For purposes of discussion, an example of the WAN is the linear broadband network 2, such as the existing HFC network, and an example of the LAN is a communication network at a subscriber premises 66. The WAN connection in the NAID 6 provides all of the functionality of a cable modem. In a preferred embodiment, the cable modem portion of the NAID 6 is DOCSIS compliant and provides DOCSIS functionality including, but not limited to, automatic negotiation, registration, encryption, and automatic assignment of IP addresses. There is a plurality of the NAIDs 6 on the linear broadband network 2. The SAID 10 communicates with one of the NAIDs 6 and provides the LAN connectivity at the subscriber premises 66 or directly at a client device. In a preferred embodiment, the subscriber access interface device 10 and the network access interface device 6 use half-duplex communications, however, full-duplex communications are also appropriate depending upon a specific application. The network access interface device 6 receives downstream information from a headend 21 and wirelessly relays the downstream information to the appropriate subscriber access interface device 10. The subscriber access interface device 10 further distributes the downstream information to an appropriate destination client device 71. Similarly, the subscriber access interface device 10 receives upstream information from an initiating client device 1 and wirelessly relays it to the network access interface device 6. The network access interface device 6 either distributes the information to another subscriber access interface device 10 supported by the same network access interface device 6 or further relays the

upstream information to the headend 21 via the linear broadband network 2. The destination client device 71 is a uni–directional or bi–directional communication device and may be, for example, a telephone, a video device, a computer, or an audio device. As one of ordinary skill in the art appreciates, therefore, the initiating and destination client devices 1, 71 may be and are typically the same device. In terms of their logical processing, however, they perform separate functions and are therefore discussed separately. The network access interface device 6 that services the subscriber access interface device 10 is that network access interface device 6 that has the "best" connection to the subscriber access interface device 6 that is being serviced. Upon occasion, however, a transient obstruction will cause a degradation of a primary communications link 102 between the subscriber access interface device 10 and the network access interface device 6. When the primary communications link 102 degrades to a predetermined threshold, the subscriber access interface device 10 searches for an alternate communications link 65 to establish with an alternate network access interface device 31.

The subscriber access interface device 10 as disclosed in the present patent application can take many forms, which will be described. Because there are so many possible variations of subscriber access interface device 10 to destination client device 71 and initiating client device 1 to network access interface device 6 embodiments, the specific embodiments described are for illustrative purposes only and do not represent all possible combinations.

The network access interface device 6 as disclosed in the present patent application is for coupling to the existing linear broadband network 2 infrastructure. An example of the linear broadband network 2 is a conventional hybrid fiber coaxial network ("HFC network") currently employed by the cable television industry for servicing its subscribers with downstream cable television. Advantageously, the network access interface device 6 is able to co–exist on the same infrastructure as a conventional wired subscriber tap 67 currently used for conventional signal distribution to a subscriber premises 66. Other examples of the linear broadband network 2 include a fiber network and a coaxial cable network that are able to reach a number of subscribers. The network access interface device 6, coupled to the linear broadband network 2, services up to sixteen subscriber premises 66 and 256 client devices 1, 71 within an approximately 300 meter radius. As a practical matter, the density of subscriber access interface devices 10 may dictate that the actual

radius of coverage be more or less than 300 meters and may vary for each subscriber access interface device 10 in the system. The actual number of subscriber premises 66and client devices 1,71 each network access interface device 6 is able to service depends upon a number of factors such as an addressing design of the network access interface device 6, a distance of the subscriber premises 66 from the network access interface device 6, a number of client devices 1, 71 within the subscriber premises66, the bandwidth of wireless channels utilized by the network access interface device 6, transmission power levels, and other factors that will occur to one of ordinary skill in the art. In one embodiment of the system according to the teachings of the present invention, the subscriber access interface device 10 may be located at the subscriber premises 66. In this embodiment, the subscriber premises 66 may be a home or a business that requires connectivity to the linear broadband network 2. The subscriber access interface device 10 can provide a minimum of the same service available using the conventional wired tap 67. A cable service provider, therefore, is able to upgrade the capability of all of its subscribers without compromising the basic service to which the subscriber had become accustomed. Additionally, the system according to the teachings of the present invention permits multiple uni--directional and bi--directional client devices 1,71 to be connected to the linear broadband network 2 via the subscriber access interface unit 10 and also permits a client device 1, 71 to establish or maintain connectivity with the linear broadband network 2 while roaming. The term "roaming" is used to describe having connectivity with the communications infrastructure while moving and without requiring physical location within the subscriber premises 66. Accordingly, a system according to the teachings of the present invention provides the minimum cable TV distribution capability while also providing other wireless communications capability on the same linear broadband network 2.

With specific reference to Figure 2 of the drawings, there is shown various possible configurations of client device 1,71 to subscriber access interface device 10. In a first embodiment a plurality of the client devices 1 are connected to the subscriber access interface device 10 via a wired connection. In this embodiment, the subscriber access interface device 10 is a peripheral device and typically has the appearance, in cable TV parlance, of a "set top box". With additional reference to Figure 3 of the drawings, the subscriber access interface device 10 receives a plurality of upstream baseband signals 3 or a plurality of upstream information

signals 11. The upstream information signals 11 are converted to the upstream baseband signal 3 format either in the initiating client device 1 or in the subscriber access interface device 10. With additional reference to Figure 4 of the drawings, the upstream baseband signal 3 format comprises a plurality of upstream data packets 24 having an address header 25. The address header 25 indicates the destination client device 71 for the upstream data packet 24 in a dotted quad IP format.

In a second embodiment of a subscriber access interface device 10 according to the teachings of the present invention, a plurality of the client devices 1 are connected to the subscriber access interface device 10 via a wireless connection. The client devices 1 modulate the upstream baseband signal 3 onto remote upstream wireless carriers to produce remote upstream modulated carriers 83. It is preferred that the remote connection between the client devices 1 and the subscriber access interface device 10 in a wireless communication embodiment follow the Home RF or Bluetooth communication protocols, but the IEEE 802.11 protocol is also an option. The remote upstream modulated carrier 83 is demodulated to reproduce the upstream data packets 24.

In a third embodiment according to the teachings of the present invention, the subscriber access interface device 10 is unitary with the client device 1. As one of ordinary skill in the art appreciates, the upstream baseband signal 3 is wired directly from the electronics in the client device 1 to the transmission electronics in the subscriber access interface device 10. The client device 1, therefore, is free to roam and communicates with whichever one of the network access interface devices 6 is local to it.

Fourth, fifth, and sixth embodiments according to the teachings of the present invention illustrate combinations of the wired and wireless connections between the client devices 1 and the subscriber access interface device 10. Embodiments five and six illustrate that the client device 1 may be a hub or router type of device with either a wired or a wireless connection. In addition the hub or router type of device may further support and be combined with other client devices 1.

Embodiments one through six as shown on Figure 2 of the drawings are intended to be illustrative and not exhaustive. Other combinations will occur to those of ordinary skill in the art and are within the scope of the present invention.

To establish upstream connectivity according to the teachings of the present invention and with reference to Figures 3 and 4 of the drawings, the initiating client device 1 generates an upstream information signal 11. The upstream information signal 11 is converted into an upstream baseband signal 3 having an Internet protocol ("IP") format using a dotted quad addressing protocol in the address header 25 followed by the upstream information data. It is not necessary that the client device 1 directly generate the upstream baseband signal 3. As represented in Figure 3 by three alternative paths for the upstream information signal 11 generated by the initiating client device 1, either the initiating client device 1, the subscriber access interface device 10, or a separate formatting device may digitize, convert, and format the upstream information into the upstream baseband signal 3.

An analog initiating client device 1, such as a legacy telephone, generates an upstream analog signal 11, such as a POTS signal. In order to convert the upstream analog signal 11 to the upstream baseband signal 3, the upstream analog signal 11 is digitized to produce a corresponding upstream digital signal. The upstream digital signal is then encoded and formatted into the upstream baseband signal 3. Alternatively, the digitizing, encoding, and formatting operations may occur within the subscriber access interface device 10, in which case the subscriber access interface device 10 receives the upstream analog signal 11 directly. The digitizing, encoding, and formatting operations may also occur in a peripheral device, which outputs the upstream baseband signal 3 for transmission to the subscriber access interface device 10. The digitizing, encoding, and formatting operations may occur within the initiating client device 1, in which case, the initiating client device 1 outputs the upstream baseband signal 3 directly to the subscriber access interface device 10. Other variations for providing the upstream baseband signal 3 to the subscriber access interface device 10 are possible and are within the scope of the present invention.

The connection from the initiating client device 1 to the subscriber access interface device 10 may be a wired connection or a wireless connection. In the case of a wireless connection between the initiating client device 1 and the subscriber access interface device 10, known communication protocols such as Home RF, Bluetooth, or IEEE 802.11 are appropriate. Specifications of the Home RF and Bluetooth communication protocols are hereby incorporated by reference. Accordingly, prior to further processing , the upstream information signal 11 is

converted into the upstream baseband signal 3 having the IP format. The upstream baseband signal 3 comprises a series of upstream data packets 24. Each upstream data packet 24 comprises a service data unit ("SDU") 48 containing upstream information to be sent to the destination client device 71 with an address header 25 in a Transmission Control Protocol/User Data Protocol ("TCP/UDP") format. The address header 25 contains one or more destination addresses, the type of data in the SDU 48 (i.e. voice, video, or data), and the total number of bytes that make up the SDU 48. As one of ordinary skill in the art appreciates, the address header 25 contains the information necessary to route the upstream data packet 24 to the appropriate destination client device 71 in the same way connectivity is established on the Internet.

As shown in Figure 2 of the drawings, the subscriber access interface device 10 is able to accept simultaneous inputs from multiple ones of the client devices 1. The subscriber access interface device 10 time domain multiplexes the upstream data packets 24, to produce a multiplexed upstream baseband signal 23. As part of the multiplexing process, the subscriber access interface device 10 may assign an even priority among all of the upstream data packets 24 or it can perform a prioritization function for each upstream data packet 24. The prioritization process provides quality of service by controlling a latency of each upstream data packet 24. As one of ordinary skill in the art will appreciate, the upstream data packets 24 from a voice or a video communications signal must be reliably and consistently delivered in real time, within the appropriate bandwidth, and cannot tolerate lost packet retransmission. The upstream data packets 24 from a data communications signal, however, may be delivered as packet bursts and are able to tolerate lost packet retransmission. Accordingly, a prioritizing process takes a data packet type into account when determining when the upstream data packet 24 is to be launched onto the multiplexed upstream baseband signal 23.

With specific reference to Figure 3 of the drawings, a central processing unit ("CPU") 60 in the subscriber access interface device 10 (hereinafter termed a "SAID CPU 60") performs the prioritization and multiplexing operations. A dedicated multiplexer performing the same functions is also appropriate. The SAID CPU 60 receives a plurality of the upstream data packets 24 and can assign a priority to each upstream data packet 24 in one of two ways.

A first prioritizing method is to assign a priority for each upstream data packet 24 based upon a source port configuration of the subscriber access interface device 10. A priority of a single packet consists of two parts, a user priority that can be either normal or high, and a maximum latency representing a maximum amount of time a packet may be delayed between sending and receiving. The maximum latency is dynamic and decrements in value as the frame waits in a queue. The priority is determined from a combination of the user priority value and the maximum latency value and is assigned a priority value from 0 to 4. Higher priority packets are transmitted first. For example, if the subscriber access interface device 10 were configured with ten ports according to embodiment 1 of Figure 2, a fraction of the ten ports is assigned as a voice port 68, a fraction is assigned as a video port 69, and a remaining fraction is assigned as a data port 70. Based upon the source port configuration, of which the subscriber access interface device 10 has a priori knowledge, the subscriber access interface device 10 generates a high priority and a minimum latency to each one of the voice ports 68 and video ports 69. Signals from each voice port 68 have the same minimum latency, which may be different than the minimum latency assigned to signals from the video ports 69. As part of the prioritization process the SAID CPU 60 in the subscriber access interface device 10 accepts a plurality of the upstream data packets 24 from the voice and video ports 68,69 according to their minimum latency and assigns the priority to each upstream data packet 24. The SAID CPU 60 then interleaves the upstream data packets 24 that are incoming from the data ports 70, assigns their priority, typically a "normal" priority and a longer minimum latency, and generates a multiplexer header 84 to reflect the assigned priorities. The first prioritizing method is most appropriate for the wired connection between the plurality of client devices 1 and the subscriber access interface device 10 and cannot be used for the wireless connection between the client devices 1 and the subscriber access interface device 10.

A second prioritizing method is to interpret the address header 25 of each upstream data packet 24. With specific reference to Figures 3 and 4 of the drawings, as part of the address header 25, each upstream data packet 24 carries an indication of the type of communications signal contained therein. The SAID CPU 60 interprets the address header 25 to determine the type of data contained in the SDU 48. Based upon the determination, the SAID CPU 60 assigns the appropriate priority and the minimum latency for each upstream data packet 24. Accordingly,

the second prioritizing method requires additional processing by the SAID CPU 60. Once the appropriate priority is assigned, data packet encapsulation appending the multiplexer header 84 and interleaving remaining data packets 24 occurs as previously described. The second prioritizing method is most appropriate for the embodiments having a wireless connection between at least one of the plurality of client devices 1 and the subscriber access interface device 10, but may also be used with the wired connection between the client devices 1 and the subscriber access interface device 10. In the embodiment wherein the client device 1 and the subscriber access interface device 10 are integral with each other, prioritization, if any, occurs using either method.

The SAID CPU 60 also calculates a cyclic redundancy check ("CRC") value 86 for the plurality of the upstream data packets 24 accepted by the SAID CPU 60. The CRC value 86 is used when the packet is received for error checking purposes. The SAID CPU 60 encapsulates the plurality of prioritized upstream data packets 24 and attaches the multiplexer header 84 at the beginning of the prioritized packets and the CRC value 86 at the end of the packet to create a multiplexed upstream baseband packet, a plurality of which comprise the multiplexed upstream baseband signal 23. The SAID CPU 60 transfers the multiplexed upstream baseband signal 23 to a subscriber access interface device Media Access Control Protocol Controller ("SAID MCU") 85 for further processing. The SAID MCU 85 encapsulates a plurality of the multiplexed upstream baseband packets and generates and appends a Media Access Control ("MAC") header 87 to produce a Payload Data Unit 88 ("PDU"). The SAID MCU 85 processes the encapsulated PDU 88 into upstream transmission fragments 89 suitable for transmission and reception over the IEEE 802.11 wireless channel. The MAC header 87 includes information regarding sequencing of the transmission fragments 89, the number of bytes in the transmission fragment and a total number of fragments in the transmission for purposes of reassembling the transmission fragments 89 after reception over the wireless channel. The upstream transmission fragments 89 are modulated onto an upstream wireless radio frequency carrier 4 by a first upstream modulator 39 to produce a first upstream modulated carrier signal 5. The first upstream modulated carrier signal 5 is a wireless communications signal and is wirelessly transmitted by the upstream transmitter 40 for reception by the network access interface device 6. Appropriate frequencies for the upstream wireless radio frequency carrier 4 may be 915MHz, 2.4GHz, or 5.8GHz depending upon the

allocated spectrum and preferably follow a wireless communications protocol that is IEEE 802.11 compliant. The IEEE 802.11 is a wireless networking standard that is specifically incorporated by reference herein. The actual frequencies used are a function of the available spectrum in a particular locality. Alternatively, a HiperLAN2 compliant process is appropriate. HiperLAN2 is also a wireless networking standard that is specifically incorporated by reference herein. The wireless link modulation and demodulation processes can employ a direct sequence spread spectrum process ("DSSS"), a frequency hopping spread spectrum process ("FSSS"), or a vector modulation process. The vector modulation processes preferably employ a quadrature phase shift keying process, but a bi–phase shift keying process, or any other modulation and demodulation process that is consistent with digital modulation techniques will work.

The first upstream modulated carrier signal 5 is received by an upstream receiver 41 in the network access interface device 6 and is demodulated by an upstream demodulator 42 in the network access interface device 6 to reproduce the data transmission fragments 89 in the network access interface device 6. A Media Access Control Processor 93 in the network access interface device 6 ("NAID MCU 93") re–constructs the transmitted encapsulated PDU 88 in the network access interface unit 6 to generate an upstream demodulated baseband signal 7. The upstream demodulated baseband signal 7 is a reconstructed version of the multiplexed upstream baseband signal 23 from which it is derived. It is not important that the upstream demodulated baseband signal 7 follow a particular format, but an IP format is preferred, and the format must preserve the prioritization, multiplexing, and destination addressing information. The general format for both the multiplexed upstream baseband signal 23 as well as the upstream demodulated baseband signal 7 should be digital and packet based. Due to the digital nature of the modulation and demodulation steps, the reconstruction process does not regenerate out of band noise that is injected into the upstream baseband signal 3 through wiring at the subscriber premises 66.

The first upstream modulated carrier signal 5 may be filtered either before transmission from the subscriber access interface device 10 or as it is received by the network access interface device 6. Filtering after reception by the network access interface device 6 is preferred. The inherent bandwidth limiting aspect of digital sampling further band limits the first upstream modulated carrier signal 5 to remove

noise products from the wireless transmission process. Accordingly, the upstream demodulated baseband signal 7 recaptures the information integrity of the multiplexed upstream baseband signal 3 and disposes of the out of band noise presented on the upstream baseband signal 3 as it is transmitted from the client device 1 through the subscriber access interface device 10 and to the network access interface device 6. In certain cases, the out of band noise on the upstream baseband signal 3 gives rise to errors in the signal recovered after the demodulation process. By using appropriate forward error correction, preferably convolution encoding and decoding, the errors in the demodulation process may be identified and corrected to further restore the information integrity of the original upstream baseband signal 3.

A Processor 90 in the network access interface device 6 (hereinafter "NAID CPU 90") interprets the address header 25 of each upstream data packet 24 to determine a specified destination address. A look up table stored in memory in the network access interface unit 6 identifies the subscriber access interface unit or units 10 that support the destination address. If the destination address points to one of the subscriber access interface devices 10 that is serviced by the same network access interface device 6, the upstream data packet 24 having the specified destination address is forwarded to the downstream path and is transmitted to the appropriate destination subscriber access interface device 10 without being forward through the headend 21. If the destination address 72 matches that of one of the subscriber access interface devices 10 that is supported by the receiving network access interface device 6, the receiving network access interface device 6 strips the upstream data packet 24 from the upstream demodulated baseband signal 7 and re-packetizes it for transmission on the downstream path to the subscriber access interface device 10 designated by the destination address. Those upstream data packets 24 having destination addresses not found in the local look up table of the network access interface device 6 remain in the upstream demodulated baseband signal 7 for transmission to the headend 21. In relaying the upstream data packet 24 to the headend 21, a second upstream modulator 43 uses the upstream demodulated baseband signal 7 to modulate an upstream linear broadband radio frequency carrier 8 to produce a second upstream modulated carrier signal 9. The second upstream modulated carrier signal 9 is in a format that is compatible with the linear broadband network 2. For example, the second upstream modulator 43 can be a DOCSIS modem to modulate and transmit the second upstream modulated carrier signal in a

DOCSIS format. The second upstream modulator 43 in the network access interface device 6 modulates the upstream linear broadband radio frequency carrier 8 and launches the second upstream modulated carrier signal 9 onto the linear broadband network 2 for transmission to the headend 21 for further distribution to the appropriate subscriber access interface device 10 and the destination device 71. Because the out of band noise in the upstream baseband signal 3 is removed, the tree and branch configuration present in the linear broadband network 2 does not present the same noise summing problems in the upstream path that are present in the conventional system. Accordingly, a method according to the teachings of the present invention is able to effectively utilize the part of the network spectrum allocated to the upstream path, on the linear broadband network, for example the 5–40MHz spectrum on an HFC network.

In order to build the local look up table in the network access interface device 10, the subscriber access interface device 6 and the network access interface device 10 perform a negotiation in which, the subscriber access interface device 6 initiates a request for negotiation. The request occurs over a separate wireless service channel initiated by the subscriber access interface device 10 when a client device configuration changes. A change occurs when a client device is added or removed from a subscriber access interface device 10. The subscriber access interface device 10 transmits information to the network access interface device 6 concerning a number of client devices supported, the information type (i.e. voice, video, or data) of each supported client device 1,71, the addresses of each client device 1,71 as well as a requested bandwidth. The network access interface device 6 accepts the information and determines how it can most best and most efficiently support the change to be made and responds to the initiating subscriber access interface device 10 a level of service it will receive. The subscriber access interface device 10 accepts the level of service and establishes communication with the network access interface device 6 thereafter. The network access interface device modifies its local look–up table according to the new information received by the subscriber access interface device 10. In certain cases, the network access interface device 6 may not be able to fully accommodate the configuration change in which case it either refuses the requested configuration change, accommodates the configuration change, but reduces the level of service for other ones of the client devices in the configuration,

or in rare cases, will reduce the level of service for one of the subscriber access interface devices that did not request a configuration modification.

With specific reference to Figure 7 of the drawings, there is shown an embodiment of a Direct Sequence Spread Spectrum ("DSSS") Network Access Interface Device ("NAID") 702 according to the teachings of the present invention. The DSSS NAID 702 is made up of five major sections: a downstream modem, an upstream modem, an access point interface, a DSSS wireless downstream block, and a DSSS wireless upstream block.

In a specific embodiment as shown in Figure 7 of the drawings, a 64/256QAM modulated downstream RF signal having a frequency range of between 54 and 860MHZ is received by a diplex filter 714 from the linear broadband network (2 in Figures 1 and 2 of the drawings). The downstream RF signal is filtered, amplified, and filtered prior to being twice down-converted in an RF tuner 716. The RF tuner 716 processes the downstream RF signal into a 36/44 MHz IF signal. A 10-bit downstream analog to digital converter 718 samples the analog IF signal converting it to a digital waveform. A quadrature IF demodulator 720 demodulates the 64/256QAM digital waveform with recovered clock and carrier timing, filters, and equalizes the data. The result of the demodulation is transmitted to a downstream forward error correction processor 722, which synchronizes, de-interleaves, decodes using a Reed-Solomon polynomial, and de-randomizes the data. The error corrected and decoded output is transmitted to a downstream processor 724 that performs DOCSIS compliant physical and Media Access Control functions. Specifically, the downstream processor 724 provides concatenation of downstream fragments and provides filtering of up to 256 destination addresses. The downstream processor 724 also provides system timing and synchronization, quality of service prioritization, bandwidth allocation, error detection, error handling, error recovery, and performs negotiation procedures with the headend (21 in Figure 1 of the drawings) for registering new NAIDs (6 in Figures 1 and 2 of the drawings). The downstream data is then sent to a system Central Processing Unit ("CPU") 726. The system CPU 726 provides 56-bit Data Encryption Standard ("DES") key information to a DES decryption processor 728 to perform privacy decryption functions. The Data Encryption Standard is hereby incorporated by reference. The system CPU 726 also provides control of the synthesizer functions for appropriate setting of the mixing and demodulation frequencies. The system

CPU 726 also contains network management protocols such as SNMP, DHCP and NAT and Internet Protocol router capability to route data packets to appropriate SAIDs 10 based upon destination addresses. System memory 734 includes a direct memory access ("DMA") controller for handling the high data rate without intervention from the system CPU 726. The DMA controller manages the transfer of the decrypted data to the system memory 734 via a system data bus 730. The DMA controller manages the transfer of the data to an IEEE 802.11 Media Access Control ("MAC") protocol controller 732 for processing. As one of ordinary skill in the art appreciates, there is a high data rate to and from the MAC protocol controller 732. Accordingly, there is a need to buffer data into and out of the MAC protocol controller 732 using a system memory 734. The system CPU 726 provides information to the MAC protocol controller 732 regarding a destination SAID 10. This triggers the MAC protocol controller 732 to send a Request to Send ("RTS") signal directly to the destination SAID 10 . When the destination SAID 10 is ready to receive data, it issues a Clear to Send ("CTS") signal. The MAC protocol controller 732 receives the CTS signal. The MAC protocol controller 732 then formats the payload data unit (92 in Figure 6 of the drawings) and appends to it a header (95 in Figure 6 of the drawings) to generate a downstream data packet (75 in Figure 6 of the drawings). The MAC protocol controller 732 transmits the downstream data packet to a digital baseband processor 738. The digital baseband processor 738 clocks and scrambles the downstream data packet and differentially encodes the downstream data packet before applying a spread spectrum modulation. The digital baseband processor 738 quadrature phase shift key modulates the packet to generate a baseband signal having I and Q components. The digital baseband processor 738 then spreads the I and Q symbols with a pseudo–random number sequence generator and sends it to a digital to analog converter 740 where the baseband signal is converted to an analog waveform. The analog waveform is transmitted to a quadrature IF modulator 742 that provides shaping and filtering and up–converts it into an IF signal. The gain controlled IF signal is further up–converted to the 2.4 to 2.5GHz band by a transmission mixer 744. The output signal from the transmission mixer 744 is filtered at 756 and power controlled at 758 to an optimized output level before it is fed into a transmit/receive diversity switch 746. An optimum power level is application dependent and differs for each network access interface unit in a system. The diversity switch 746 connects the output

signal to a transmit antenna 748 for transmission by the NAID 6 and subsequent reception by the SAID 10. AMAC switch control line 736 controls the position of the diversity switch 746 to connect the transmit antenna 748 for delivery of the request to send signal. The MAC switch control line 736 then switches the position of the diversity switch 746 to connect the receive antenna 750 for reception of the clear to send signal. As described and shown in Figure 7, the NAID 6 may comprise separate transmit and receive antennas 748, 750 respectively with a double pole/double throw diversity switch 746. Alternatively, the NAID may comprise a single transmit/receive antenna 752 for use with a single pole/double throw diversity switch 754.

With further reference to Figure 7 of the drawings, an upstream wireless signal having a frequency in the 2.4 to 2.5GHz range is received by the receive antenna 750 and filtered with a dielectric filter (not shown). The diversity switch 746 is in a position connecting the receive antenna 750 to wireless receive circuitry in the NAID 6. A low noise amplifier 760 sets the received noise figure and appropriate signal gains of the received signal. After filtering at 762, the signal is down–converted in mixer 764 to a 280MHz IF signal. The IF signal is bandpass filtered at 766 and is transmitted to quadrature demodulator 768. The quadrature demodulator 768 comprises a limiting amplifier, baseband demodulator, and baseband low pass filters. The resulting I and Q quadrature signals are converted to digital waveforms in the analog to digital converter 770 and transmitted to the digital baseband processor 738. The digital baseband processor 738 correlates the pseudo–random number spreading to remove it and to recover the differential quadrature phase shift keying data. The digital baseband processor detects, identifies, and locks onto the signal and uncovers the symbol timing phase and frequency. The digital baseband processor 738 uses the symbol timing phase and frequency to initialize a tracking loop for data acquisition. When the digital baseband processor 738 begins successful tracking of the demodulated data, it de–scrambles the direct spread data to prepare the upstream transmission fragments (89 in Figure 4 of the drawings) for processing in the IEEE 802.11 MAC protocol controller 732. The MAC protocol controller 732 provides the concatenation function to recover the upstream PDU (88 in Figure 4 of the drawings). Each encapsulated PDU comprises the MAC header (87 in Figure 4) containing a preamble and a start frame delimiter, the data, and the CRC value (86 in Figure 4). The MAC protocol controller 732 processes and

interprets the MAC header and start frame delimiter, determines a mode and length of incoming PDU, and checks the CRC value. If the CRC value indicates that the data is corrupted, the MAC protocol controller 732 discards the current packet and issues a retransmission request. If the CRC value indicates that the data is acceptable, the MAC protocol controller 732 further processes the packet to strip off the MAC header 87 and launches the packet onto the system data bus 730 for delivery to the system CPU 726 via the system memory 734. The system CPU 726 supplies a 56-bit DES key to a DES encrypt block 772, which encrypts each packet received. Each encrypted packet is sent to an upstream processor 774 that handles elements of time synchronization with the headend (21 in Figure 1), bandwidth request negotiation, and contention resolution. The packets are then Reed-Solomon encoded for forward error correction in an upstream forward error correction block 776. The upstream forward error correction block also randomizes, appends a preambleto a beginning of the packet, maps the QPSK/QAM modulation symbols, and pre-equalizes the transmission signal. At this point in the process, an output of the upstream forward error correction block 776 is a digital shaped pulse waveform. The digital waveform is then upconverted to an IF central frequency by a quadrature IF modulator 778 generating an output data burst containing data at a variable symbol rate in either a QPSK or 16-QAM format. A 10-bit digital to analog converter 780 converts the output data burst signal to a 5-65MHz analog waveform. The analog waveform is filtered at 782 and amplified by an automatic gain controlled power amplifier 784. The amplified signal is filtered again at 786 before being launched onto the linear broadband network 2 through the diplexer filter 714. Power is delivered to the NAID via the coaxial connection to the linear broadband network (2 in Figure 1 of the drawings). The AC power signal, at for example 60Hz or 50Hz, is superimposed onto the communications signal. The diplexer filter 714 separates the low frequency power signal from the higher frequency communications signal. The low frequency power signal is sent to power supply 788 where the AC power is converted to DC using known techniques and then distributed throughout the NAID over power busses 790.

With specific reference to Figure 8 of the drawings, there is shown a direct spread spectrum subscriber access interface unit block diagram 802 according to the teachings of the present invention in which an IEEE 802.11 wireless DSSS signal is received by receive antenna 804. The primary downstream function of the

subscriber access interface unit 802 is delivery of the appropriate signal to a designated destination client device 830(71 in Figure 2). With a diversity switch 806 in a receive position, the received signal is transmitted to a low noise amplifier 808, which sets the received noise figure and appropriate signal gains of the received signal. After filtering at 810, the signal is down–converted in mixer 812 to a 280MHz IF signal. The IF signal is band–pass filtered at 814 and is transmitted to quadrature demodulator 816. The quadrature demodulator 816 comprises a limiting amplifier, baseband demodulator, and baseband low pass filters. The resulting I and Q quadrature signals are converted to digital waveforms in analog to digital converter 818 and transmitted to the digital baseband processor 820. The digital baseband processor 820 correlates the pseudo–random number spreading to remove it and to recover the differential quadrature phase shift keying data. The digital baseband processor 820 detects, identifies, and locks onto the signal and uncovers the symbol timing phase and frequency. The digital baseband processor 820 uses the symbol timing phase and frequency to initialize a tracking loop for data acquisition. When the digital baseband processor 820 begins successful tracking of the demodulated data, it de–scrambles the direct spread data to prepare the downstream transmission fragments (94 in Figure 6 of the drawings) for launch onto a system data bus 822 and processing in IEEE 802.11 MAC protocol controller 824. The MAC protocol controller 824 provides the concatenation function to recover the downstream data packets (75 in Figure 6 of the drawings). Each packet comprises the MAC header (87 in Figure 6) containing a preamble and a start frame delimiter, the data, and the CRC value (86 in Figure 4). The MAC protocol controller 824 processes and interprets the MAC header (87 in Figure 6) and start frame delimiter, determines a mode and length of the incoming packet, and checks the CRC value (86 in Figure 6). If the CRC value indicates that the data is corrupted, the MAC protocol controller 824 discards the current packet and issues a retransmission request. If the CRC value indicates that the data is acceptable, the MAC protocol controller 824 further processes the packet to strip off the MAC header and launches the packet onto the system data bus 822 for delivery to a system CPU 826 via the system memory 828. A system interface that comprises the digital baseband processor 820, the system data bus 822, the MAC protocol controller 824, the system memory 828, and the system CPU 826, is interrupt driven to support the high data transmission speeds. The system memory 828 comprises a direct memory access controller to

handle actual data transfer to and from the system memory 828 without intervention from the system CPU 826 or MAC protocol controller 824. The MAC protocol controller 824 and the system CPU 826 are notified via hardware interrupt when data is ready for processing and when a memory receive buffer is full. The system CPU 826 receives and interprets the address header (25 if Figure 6) of each data packet and de–multiplexes each packet for transmission to the designated destination client devices 830 (71 in Figure 2).

With further reference to Figure 8 of the drawings, a primary upstream function of the subscriber access interface unit is gathering, encoding, and multiplexing the signals from a plurality of initiating client devices 832 (1 in Figure 2 of the drawings) onto a single signal for wireless transmission to the NAID (6 in Figure 1 of the drawings). Figure 8 of the drawings represents a DSSS embodiment of the SAID in which a plurality of initiating client devices 832 generate upstream signals. Each upstream signal is encoded into an upstream baseband signal (3 in Figure 3 of the drawings). The system CPU 826 receives each upstream baseband signal 3, prioritizes, multiplexes, and generates and appends the address header (25 in Figure 4 of the drawings) and the multiplexer header (84 in Figure 4 of the drawings) onto each packet before launching onto the system data bus 822 for storage into the system memory 828. The MAC protocol controller 824 receives the plurality of packets stored in the system memory 828 and encapsulates and then fragments the packet for transmission over the IEEE 802.11 wireless link. The digital baseband processor 820 clocks and scrambles the downstream transmission fragments (89 in Figure 4 of the drawings) and differentially encodes the downstream transmission fragments before applying a spread spectrum modulation. The digital baseband processor 820 quadrature phase shift key modulates each fragment to generate a baseband signal having I and Q components. The digital baseband processor 820 then spreads the I and Q symbols with a pseudo–random number sequence generator and sends it to a digital to analog converter 834 where the baseband signal is converted to an analog waveform. The analog waveform is transmitted to a quadrature IF modulator 836 that provides shaping and filtering and up–converts it into an IF signal. The gain controlled IF signal is further up–converted to the 2.4 to 2.5GHz band by a transmission mixer 838. The output signal from the transmission mixer 838 is filtered at 840 and power controlled at 842 to an optimized output level before it is fed into the transmit/receive diversity switch 806.

An optimum power level is application dependent and differs for each subscriber access interface unit in a system. The diversity switch 806 connects the output signal to a transmit antenna 844 for wireless transmission by the NAID and subsequent reception by the SAID. AMAC switch control line 846 controls the position of the diversity switch 806 to connect the transmit antenna 844 for delivery of the request to send signal. The MAC switch control line 846 then switches the position of the diversity switch 806 to connect the receive antenna 804 for reception of the clear to send signal. As described and shown in Figure 8, the SAID 6 may comprise separate transmit and receive antennas 844, 804 respectively with a double pole/double throw diversity switch 806. Alternatively, the SAID may comprise a single transmit/receive antenna 848 for use with a single pole/double throw diversity switch 850. The MAC switch control line 846 controls the diversity switch 806, 850 similarly in both alternatives.

With specific reference to Figure 9 of the drawings, there is shown a frequency hopping spread spectrum ("FHSS") embodiment of the network access interface device 902 according to the teachings of the present invention in which all elements are the same as described in the DSSS embodiment of the NAID 702 except for the transmission and reception circuitry disposed downstream and upstream, respectively, from a digital baseband processor 904. Accordingly, for purposes of eliminating duplication, it is the differing portion of the FHSS embodiment that is discussed in this paragraph. In the downstream process, the digital baseband processor 904 convolution encodes the data packets received from a MAC protocol controller 906 via a system data bus 908 into I and Q digital data channels. The digital baseband processor 904 sends the I and Q digital data channels to downstream analog baseband processor 910. The downstream analog baseband processor 910 filters, differentially encodes, and converts the digital signals to filtered and encoded analog equivalent signals. Accordingly, the data is QPSK modulated and comprises a baseband quadrature signal having I and Q components. The resulting signal is up-converted to an intermediate frequency in IF mixer 912. An input to the mixer 912 is an output of a frequency hopping generator 916. The generator 916 comprises two voltage controlled oscillators 918, 920 ("VCO") in a phase locked loop and a frequency hopping sequence controller 917. While one VCO 918 is operating as input to the IF mixer 912, the other VCO 920 is slewing to a new frequency. The MAC protocol controller 906 directs the frequency hopping

sequence controller 917 as to the appropriate frequency to generate. Switch 922 toggles between the two VCOs 918, 920 as the frequencies slew between the various values. The IF signal is then filtered at 924 and up–converted again in mixer 926 to a 2.4GHz to 2.5GHz RF signal. The up–converted RF signal is then filtered at 928 to remove the IF image frequency. The filtered and up–converted signal is then amplified in power amplifier 930. Triple pole, double throw diversity switch 932 is in a transmit position and directs the amplified signal to transmit antenna 934 for wireless transmission to a receiving FHSS SAID. Alternatively, the FHSS NAID embodiment may use a single transmit/receive antenna 936 and a single pole/double throw diversity switch 938.

With further reference to Figure 9 of the drawings, the FHSS NAID receives an IEEE 802.11 wireless signal at receive antenna 940. With the diversity switch 932 or 938 in a receive position, the received signal is filtered at dielectric filter 942. Low noise amplifier 944 sets the received noise figure and appropriate signal gain. The signal is then down–converted to a 280MHz IF signal in RF mixer 946 and then filtered at 948 and further down–converted to reconcile the frequency hopping sequence in mixer 950. The frequency input to the mixer 950 comprises the output of the frequency hopping generator 916 as directed by the MAC protocol controller 906 and controlled by the frequency hopping sequence controller 917. The resulting IF signal is input into a downstream analog baseband processor 952. The analog baseband processor 952 down–converts the signal to baseband producing I and Q signals and digitizes the analog signals in a 10–bit analog to digital converter for transmission to the digital baseband processor 904. The digital baseband processor 904 performs a complex frequency rotation to adjust for any frequency offset and phase error between the transmitter in the SAID and the receiver in the NAID. The digital baseband processor 904 then provides symbol timing and carrier frequency acquisition and tracking. The digital baseband processor 904 also provides automatic gain control on the demodulated baseband signal and a decision threshold comparison of the I and Q channel against an appropriate reference level. The pair of I and Q soft decision signals isthen sent to a Viterbi Decoder portion of the digital baseband processor 904. The digital baseband processor 904 also determines the synchronization boundary of the QPSK symbols, performs the forward error correction decoding process, and recovers the data transmission fragments for interpretation and processing by the MAC protocol controller 906. Each

transmission fragments has a preamble and header containing a start frame delimiter, data, and a CRC value. The Mac protocol controller 906 processes the start frame delimiter and the header, checks the CRC value, and determines the mode and the length of the incoming message. If the CRC value indicates that the data is corrupt, the MAC protocol controller 906 issues a retransmission requestfor delivery to the SAID. If the CRC value indicates that the data is not corrupt, MAC protocol controller 906 recovers the packet and sends it to a system CPU 956 via the system data bus 908 and system memory 958.

With specific reference to Figure 10 of the drawings, there is shown a FHSS embodiment of the SAID according to the teachings of the present invention in which all elements are the same as described in the DSSS embodiment of the SAID 802 except for the transmission and reception circuitry disposed downstream and upstream, respectively, from a digital baseband processor 1002 in Figure 10 (820 in Figure 8). Accordingly, for purposes of eliminating duplication, it is the different portion of the FHSS embodiment that is discussed in this paragraph. In the upstream process, the digital baseband processor 1002 convolution encodes the data packets received from a MAC protocol controller 1004 via a system data bus 1006 into I and Q digital data channels. The digital baseband processor 1002 sends the I and Q digital data channels to upstream analog baseband processor 1008. The upstream analog baseband processor 1008 filters, differentially encodes, and converts the digital signals to filtered and encoded analog equivalent signals. Accordingly, the data is QPSK modulated and comprises a baseband quadrature signal having I and Q components. The resulting signal is up–converted to an intermediate frequency in IF mixer 1010. An input to the IF mixer 1010 is an output of a frequency hopping generator 1012. The frequency hopping generator 1012 comprises two voltage controlled oscillators 1014, 1016 ("VCO") in a phase locked loop and a frequency hopping sequence controller 1018. While one VCO 1014 is operating as input to the IF mixer 1010, the other VCO 1016 is slewing to a new frequency. The MAC protocol controller 1004 directs the frequency hopping sequence controller 1018 as to the appropriate frequency to generate. Switch 1020 toggles between the two VCOs 1014, 1016 as the frequencies slew between the various values. The IF signal is then filtered at 1022 and up–converted again in mixer 1024 to a 2.4GHz to 2.5GHz RF signal. The up–converted RF signal is then filtered at 1026 to remove the IF image frequency. The filtered and up–converted signal is then amplified in

power amplifier 1028. Triple pole, double throw diversity switch 1030 is in a transmit position and directs the amplified signal to transmit antenna 1032 for wireless transmission to a receiving FHSS NAID. Alternatively, the FHSS NAID embodiment may use a single transmit/receive antenna 1034 and a single pole/double throw diversity switch 1036 instead of separate transmit and receive antennas 1032 and 1038 respectively and diversity switch 1030.

With further reference to Figure 10 of the drawings, for downstream operation, the FHSS SAID receives an IEEE 802.11 wireless signal at the receive antenna 1038. With the diversity switch 1030 or 1036 in a receive position, the received signal is filtered at dielectric filter 1040. Low noise amplifier 1042 sets the received noise figure and appropriate signal gain. The signal is then down-converted to a 280MHz IF signal in RF mixer 1044 and then filtered at 1046 and further down-converted to reconcile the frequency hopping sequence in mixer 1048. The frequency input to the mixer 1048 comprises the output of the frequency hopping generator 1012 as directed by the MAC protocol controller 1004 and controlled by the frequency hopping sequence controller 1018. The resulting IF signal is input into a downstream analog baseband processor 1050. The downstream analog baseband processor 1050 down-converts the signal to baseband, producing I and Q signals, and digitizes the analog signals in a 10-bit analog to digital converter for transmission to the digital baseband processor 1002. The digital baseband processor 1002 performs a complex frequency rotation to adjust for any frequency offset and phase error between the transmitter in the NAID and the receiver in the SAID. The digital baseband processor 1002 then provides symbol timing and carrier frequency acquisition and tracking. The digital baseband processor 1002 also provides automatic gain control on the demodulated baseband signal and a decision threshold comparison of the I and Q channel against an appropriate reference level. The pair of I and Q soft decision signals isthen sent to a Viterbi Decoder portion of the digital baseband processor 1002. The digital baseband processor 1002 also determines the synchronization boundary of the QPSK symbols, performs the forward error correction decoding process, and recovers the data transmission fragments for interpretation and processing by the MAC protocol controller 1004. Each transmission fragment has a preamble and header containing a start frame delimiter, data, and a CRC value. The Mac protocol controller 1004 processes the start frame delimiter and the header, checks the CRC value, and determines the

mode and the length of the incoming message. If the CRC value indicates that the data is corrupt, the MAC protocol controller 1004 issues a retransmission request for receipt by the NAID. If the CRC value indicates that the data is not corrupt, MAC protocol controller 1004 recovers the packet and sends it to a system CPU 1054 via the system data bus 1006 and system memory 1056.

With specific reference to Figure 11 of the drawings, a specific beneficial application that takes advantage of the upstream communication capabilities of a system according to the teachings of the present invention is a polling system, whereby a polling request 1102 is received by the network access interface device 6 to take a measurement for immediate reporting or retrieval at a later time. The polling request 1102 may be generated at the network access interface device 6 upon initiation from the network access interface device 6 itself, another set-top box in the subscriber premises 66, a downstream client device 1, the headend 21, or from a device external to the linear broadband network 2 and contains information specifying the content of the request as well as the destination of a response to the request. Upon the receiving the polling request 1102, the network access interface device 10 generates a polling signal and transmits it to the subscriber access interface device 10 in step 1104. The subscriber access interface device 10 responds to the polling signal by generating a response upstream baseband signal 1106. The response upstream baseband signal contains the data gathered in response to the polling request 1102 as well as an address of an intended destination of the data. A data format of the response upstream baseband signal is the same as that of the upstream baseband signal 3 and is, therefore, processed similarly. The subscriber access interface device 10 modulates and transmits the response upstream baseband signal to the network access interface device 6. The network access interface device 6 in step 1108 determines whether the data contained in the response upstream baseband signal is to be stored in the network access interface device 6 for later retrieval, whether the response upstream baseband signal is to be forwarded directly to the headend 21, or whether the data is to be sent to one of the destination client devices 71. If the data contained in the response upstream baseband signal is to be stored, the network access interface device 6 receives the response upstream baseband signal, converts it to data, and stores the data in memory for later retrieval shown in step 1110. If the data contained in the response upstream baseband signal is to be transmitted to the headend 21, the network access interface device 6 receives

the response upstream baseband signal, demodulates, re–modulates, and transmits it to the headend 21 along with the upstream data shown in step 1112. If the data contained in the response upstream baseband signal is to be transmitted to one of the destination client devices 71, the network access interface device 6 receives the response upstream baseband signal, demodulates the upstream baseband signal and forwards it to the destination client device 71 shown in step 1114. Stored data may be retrieved at any time upon a request from the headend 21 or any one of the client devices 1. The polling capability has application in areas of network maintenance and meter reading, for example.

The network access interface device 6 supports a plurality of the subscriber access interface devices 6. Accordingly, the network access interface device 6 performs an arbitration process whereby the network access interface device 6 controls the timing of the receipt of transmissions from each of the subscriber access interface devices 10 supported by the network access interface device 6. Preferably, the arbitration function follows the process as defined in the IEEE 802.11 standard, the contents of which are specifically incorporated by reference herein, in which the subscriber access interface device 6 transmits a signal to the network access interface device 10 requesting access to a transmission channel. The network access interface device 10 responds with a channel clear after which the subscriber access interface device 10 transmits for a certain period of time. The transmission is acknowledged with an indication of whether the upstream data transmission fragments 89 were successfully received or not, and the process repeats for other subscriber access interface devices 10. As one of ordinary skill in the art appreciates, the network access interface device 6 may be built to accept a plurality of the first upstream modulated carrier signals 5 in order to increase upstream bandwidth between the network access interface device 6 and the subscriber access interface device 10.

Under certain circumstances, it is possible that a transient obstruction can disturb the communications link between the subscriber access interface device 10 and its associated network access interface device 6 or that a network access interface device 6 is utilizing all of its capacity for communications operations. Each subscriber access interface device 10 has a predetermined threshold against which it measures whether an existing communication link is adequate. The predetermined threshold may be a measurement of bit error rate, latency of upstream data packet transmission, signal strength, or a combination thereof. With reference

to Figure 1 of the drawings, In the event that the communications link falls below the predetermined threshold, the subscriber access interface device 10 seeks and establishes an alternate communications link 65 with an alternate network access interface device 31.

The downstream path does not share the same noise problems with the upstream path. A system according to the teachings of the present invention, however, provides the opportunity to advantageously utilize the same infrastructure for the downstream path. Additionally, the downstream path permits the cable television service providers to offer bi–directional communication services traditionally handled by the telephone companies. With reference to Figures 5 and 6 of the drawings, the downstream process comprises the network access interface device 6 accepting at least one first downstream modulated carrier signal 34 from the linear broadband network 2. A first downstream receiver 54 in the network access interface device 6 receives the first downstream modulated carrier signal 34 and transmits it to a DOCSIS compliant first downstream demodulator 55. The first downstream demodulator 55 demodulates the first downstream modulated carrier signal 34 to produce a first downstream baseband signal 35 comprising a series of downstream data packets 75. Each downstream data packet 75 has a physical address header 91, an address header 25 in an IP dotted quad format, a downstream data payload 92, and the CRC byte 86. The NAID CPU 90interprets the physical address header 91 of each of the downstream data packets 75. If the physical address points to one or more of the subscriber access interface units 10 supported by the network access interface unit 6, the NAID CPU 90 uses a downstream look–up table to re–map a physical address contained in the physical address header 91 into a corresponding logical subscriber access interface device address header 95 for further downstream processing. If the physical address does not match, the packet is not further processed. The downstream data packets 75 destined for local subscriber access interface devices 10 are further processed as part of the first downstream baseband signal 35. The NAID MCU 93 fragments the packets in the first downstream baseband signal 35 to produce downstream transmission fragments 94 and appends the MAC header 87 for purposes of sequencing and reassembly of the downstream transmission fragments 94 as they are received by the subscriber access interface device 10.

A downstream modulator 56 modulates the first downstream baseband signal 35 onto at least one downstream wireless radio frequency carrier 36 to produce at least one second downstream modulated carrier signal 37. The network access interface device 6 may also encode each downstream modulated carrier signal 37 with forward error correction encoding prior to transmission. In a preferred embodiment, convolution encoding is used for forward error correction. A downstream transmitter 57 transmits the at least one second downstream modulated carrier signal 37 wirelessly to the subscriber access interface device 10. The second downstream receiver 58 in the subscriber access interface device 10 receives the second downstream modulated carrier signal 37 and transmits it to a second downstream demodulator 59. The second downstream demodulator 59 demodulates the at least one second downstream modulated carrier signal 37 to produce at least one second downstream baseband signal 38. The SAID MCU 85 interprets the downstream MAC header 87 and reassembles the downstream transmission fragments 94 to reproduce the downstream data packet 75. If forward error correction has been used, the SAID CPU 60 also decodes each downstream data packet baseband signal 38 prior to transmission to the destination client device 71. The second downstream baseband signal 38 may, but need not, have the same format as the first downstream baseband signal 35. The information integrity, however, is preserved. It is preferred that the first downstream baseband signal 35 uses an IP format. The second downstream baseband signal 38 is a packetized signal, which is processed by the SAID CPU 60 to interpret a logical subscriber access interface device header 95 of each downstream data packet 75. The logical subscriber access interface device header 95 contains de-multiplexing information for purposes of providing quality of service to higher priority data such as voice and video. The SAID CPU 60 interprets the header and de-multiplexes the downstream data packets 75 according to their priorities. The address header 25 indicates the IP address of the destination device 71 to which the data is to be forwarded. The SAID CPU 60 interprets the address header and transmits each downstream data packet 75 to the appropriate destination client device 71.

In the case of a wired subscriber access interface device 10, the SAID CPU 60 directs each downstream data packet 75 to an appropriate data port according to a hardware configuration, each port having its own unique IP address. In the case of a wireless connection between the subscriber access interface device 10 and the

destination client device 71, the second downstream baseband signal 38 is modulated onto at least one second downstream wireless radio frequency carrier to produce a third downstream modulated carrier signal which is transmitted wirelessly. Each of the destination client devices 71 receives the third downstream modulated carrier signal and demodulates it to a third downstream baseband signal re-producing the information contained in the downstream data packets 75. Each destination client device 71 interprets the address header 25 of each downstream data packet 75 searching for a single match to the destination address. If the destination address matches the address of the destination client device 71, the destination client device 71 accepts, decodes, and presents the downstream data packet 75 containing the matching destination address. The downstream data packets 75 having destination addresses that do not match are not presented by the destination client device 71 and are discarded after interpreting the non-matching destination address. The present invention has been described by way of example. Modifications and variations to the teachings in the present disclosure are possible without departing from the scope of the appended claims.

Claims

1. A method of communicating information from a client device to a linear broadband network having substantially linear and broadband frequency characteristics comprising the steps of:

generating an upstream baseband signal, said upstream baseband signal having a predefined format,

modulating the upstream baseband signal onto at least one upstream wireless radio frequency carrier to generate at least one first upstream modulated carrier signal,

transmitting said at least one first upstream modulated carrier signal wirelessly,

receiving said at least one first upstream modulated carrier signal at a network access interface device coupled to the linear broadband network,

demodulating said at least one first upstream modulated carrier signal to produce an upstream demodulated baseband signal, and

modulating said upstream demodulated baseband signal onto at least one upstream linear broadband radio frequency carrier to produce at least one second upstream modulated carrier signal having a signal format compatible with the linear broadband network.

2. A method of communicating as recited in claim 1 and further comprising the step of processing said upstream demodulated baseband signal to substantially reconstruct said upstream baseband signal to said predefined format.

3. A method of communicating as recited in claim 1 and further comprising the step of processing said upstream demodulated baseband signal to substantially reconstruct said upstream baseband signal to a format different from said predefined format.

4. A method of communicating as recited in claim 1 wherein the step of processing further comprises the step of correcting errors present in said upstream

demodulated baseband signal to restore information integrity of said upstream baseband signal.

5. A method of communicating as recited in claim 1 and further comprising the step of filtering said at least one first upstream modulated carrier signal.

6. A method of communicating as recited in claim 5 wherein said step of filtering occurs prior to the step of transmitting said at least one first upstream modulated carrier signal.

7. A method of communicating as recited in claim 5 wherein said step of filtering occurs after the step of receiving said at least one first upstream modulated carrier signal.

8. A method of communicating as recited in claim 1 wherein said linear broadband network is a hybrid fiber coaxial network.

9. A method of communicating as recited in claim 1 wherein said linear broadband network is a linear fiber network.

10. A method of communicating as recited in claim 1 wherein said linear broadband network is a linear coaxial network.

11. A method of communicating as recited in claim 1 wherein said linear broadband network is a cable TV network.

12. A method of communicating as recited in claim 1, the step of generating and modulating occurring in said client device.

13. A method of communicating as recited in claim 1, the client device communicating with a subscriber access interface device, the subscriber access interface device performing the steps of modulating and transmitting.

14. A method of communicating as recited in claim 13, said client device communicating with said subscriber access device over a wireless connection.

15. A method of communicating as recited in claim 13, the client device communicating with said subscriber access device over a wired connection.

16. A method of communicating as recited in claim 1 wherein said step of generating said upstream baseband signal further comprises the steps of generating an upstream analog signal, digitizing said upstream analog signal to produce a corresponding upstream digital signal, and converting said upstream digital signal to said upstream baseband signal.

17. A method of communicating as recited in claim 16, wherein the step of generating said upstream analog signal comprises generating a telephone signal.

18. A method of communicating as recited in claim 17, wherein the step of generating said telephone signal comprises generating a voice signal.

19. A method of communicating as recited in claim 17, wherein the step of generating said telephone signal comprises generating a telephonic fax signal.

20. A method of communicating as recited in claim 17, wherein the step of generating said telephone signal comprises generating a telephonic data modem signal.

21. A method of communicating as recited in claim 16, wherein the step of generating said analog signal comprises generating a video signal.

22. A method of communicating as recited in claim 16, wherein the step of generating said analog signal comprises generating an audio signal.

23. A method of communicating as recited in claim 1, the step of transmitting said at least one first upstream modulated carrier signal wirelessly comprises transmitting said at least one first upstream modulated carrier signal wirelessly from each one of a plurality of subscriber access interface devices.

24. A method of communicating as recited in claim 23, the step of transmitting said at least one upstream modulated carrier signal wirelessly from each one of said plurality of subscriber access interface devices further comprising the step of arbitrating for access to said network access interface device from said plurality of subscriber access interface devices.

25. A method of communicating as recited in claim 1, the steps of modulating and demodulating said upstream baseband signal onto said at least one upstream wireless radio frequency carrier using a direct sequence spread spectrum process.

26. A method of communicating as recited in claim 1, the steps of modulating and demodulating said upstream baseband signal onto said at least one upstream wireless radio frequency carrier using a frequency hopping spread spectrum process.

27. A method of communicating as recited in claim 1, the steps of modulating and demodulating said upstream baseband signal onto said at least one upstream wireless radio frequency carrier using a vector modulation process.

28. A method of communicating as recited in claim 27, the steps of modulating and demodulating said upstream baseband signal onto said at least one upstream wireless radio frequency carrier using a quadrature phase shift keying process.

29. A method of communicating as recited in claim 27, the steps of modulating and demodulating said upstream baseband signal onto said at least one upstream wireless radio frequency carrier using a bi–phase shift keying process.

30. A method of communicating as recited in claim 1, the steps of modulating and demodulating said upstream baseband signal onto said at least one upstream wireless radio frequency carrier using an IEEE 802.11 compliant process.

31. A method of communicating as recited in claim 1, the steps of modulating and demodulating said upstream baseband signal onto said at least one upstream wireless radio frequency carrier using a HiperLAN2 compliant process.

32. A method of communicating as recited in claim 1 and further comprising the steps of:

generating a polling signal at said network access interface device,

receiving said polling signal at a subscriber access interface device,

responding to said polling signal by generating a response upstream baseband signal,

transmitting said response upstream baseband signal, and

receiving said response upstream baseband signal in said network access interface device.

33. A method of communicating as recited in claim 32 and further comprising the step of forwarding said response upstream baseband signal to a headend on said linear broadband network.

34. A method of communicating as recited in claim 32 and further comprising the step of storing said response upstream baseband signal for retrieval of said response upstream baseband signal upon request.

35. A method of communicating as recited in claim 1, the step of generating an upstream baseband signal further comprising the steps of generating a communication signal in each one of a plurality of said client devices, converting each said communication signal into a respective one of said upstream baseband signals, time division multiplexing each said upstream baseband signal in a subscriber access interface device to produce a multiplexed upstream baseband signal and the step of modulating further comprising the step of modulating said multiplexed baseband signal onto said at least one upstream wireless radio frequency carrier.

36. A method of communicating as recited in claim 35, wherein the step of converting said communication signal further comprises the step of coupling said upstream baseband signal to said subscriber access interface through a wired connection.

37. A method of communicating as recited in claim 35, wherein the step of converting said communication signal further comprises the step of coupling said upstream baseband signal to said subscriber access interface device through a wireless connection.

38. A method of communicating as recited in claim 35, each upstream baseband signal comprising a plurality of upstream data packets, the method further comprising the steps of prioritizing a launch of each upstream data packet onto said

multiplexed upstream baseband signal to control latency of each individual upstream data packet.

39. A method of communicating as recited in claim 38 and further comprising the steps of interpreting a header of each upstream data packet, assigning a priority to each upstream data packet, producing an assigned priority for each upstream data packet, and multiplexing each upstream data packet according to said assigned priority.

40. A method of communicating as recited in claim 38 and further comprising the steps of assigning a priority to each upstream data packet based upon a source port configuration, producing an assigned priority for each upstream data packet, and multiplexing each upstream data packet according to said assigned priority.

41. A method of communicating as recited in claim 1 and further comprising the step of:

controlling a timing of the step of receiving said first upstream modulated carrier signal.

42. A method of communicating as recited in claim 1 and further comprising the step of filtering a spectral range of all of said at least one upstream modulated carrier signal.

43. A method of communicating as recited in claim 1 and further comprising the step of establishing a communications link with an alternate network access interface device in the event of a communication link degradation below a predetermined threshold with said network access interface device.

44. A method of communicating bi-directional information between a client device and a linear broadband network having substantially linear and broadband frequency characteristics comprising the steps of:

generating an upstream baseband signal, said upstream baseband signal having a predefined format,

modulating said upstream baseband signal onto at least one upstream wireless radio frequency carrier to generate at least one first upstream modulated carrier signal,

transmitting said at least one fist upstream modulated carrier signal wirelessly,

receiving said at least one first upstream modulated carrier signal at a network access interface device coupled to the linear broadband network,

demodulating said at least one first upstream modulated carrier signal to produce an upstream demodulated baseband signal,

modulating said upstream demodulated baseband signal onto at least one upstream linear broadband radio frequency carrier to produce at least one second upstream modulated carrier signal having a signal format compatible with the linear broadband network,

receiving at least one downstream linear broadband network radio frequency carrier signal comprising a first downstream modulated carrier signal from the linear broadband network,

demodulating said at least one first downstream modulated carrier signal to produce at least one first downstream baseband signal having a predefined format,

modulating said at least one first downstream baseband signal onto at least one downstream wireless radio frequency carrier to generate at least one second modulated downstream carrier signal,

transmitting said at least one second modulated downstream carrier signal wirelessly,

receiving said at least one second modulated downstream carrier signal,

demodulating said at least one second modulated downstream carrier signal to produce at least one second downstream baseband signal,

transmitting said at least one second downstream baseband signal.

45. A method of communicating as recited in claim 44 and further comprising the step of processing said upstream demodulated baseband signal to substantially reconstruct said upstream baseband signal to said predefined format.

46. A method of communicating as recited in claim 44 and further comprising the step of processing said upstream demodulated baseband signal to substantially reconstruct said upstream baseband signal to a format different from said predefined format.

47. A method of communicating as recited in claim 44 wherein the step of processing further comprises the step of correcting errors present in said upstream demodulated baseband signal to restore information integrity to correspond with said upstream baseband signal.

48. A method of communicating as recited in claim 44 and further comprising the step of filtering said at least one upstream modulated carrier signal.

49. A method of communicating as recited in claim 48 wherein said step of filtering occurs prior to the step of transmitting said at least one first upstream modulated carrier signal.

50. A method of communicating as recited in claim 48 wherein said step of filtering occurs after the step of receiving said at least one first upstream modulated carrier signal.

51. A method of communicating as recited in claim 44 wherein said linear broadband network is a hybrid fiber coaxial network.

52. A method of communicating as recited in claim 44 wherein said linear broadband network is a linear fiber network.

53. A method of communicating as recited in claim 44 wherein said linear broadband network is a linear coaxial network.

54. A method of communicating as recited in claim 44 wherein said linear broadband network is a cable TV network.

55. A method of communicating as recited in claim 44, the step of generating and modulating occurring in said client device.

56. A method of communicating as recited in claim 44, the client device communicating with a subscriber access interface device, the subscriber access interface device performing the steps of modulating and transmitting.

57. A method of communicating as recited in claim 56, said client device communicating with said subscriber access interface device over a wireless connection.

58. A method of communicating as recited in claim 57, said client device communicating with said subscriber access interface device over a wired connection.

59. A method of communicating as recited in claim 44 wherein said step of generating said upstream baseband signal further comprises the steps of generating an upstream analog signal, digitizing said upstream analog signal to produce a corresponding upstream digital signal, and converting said upstream digital signal to said upstream baseband signal.

60. A method of communicating as recited in claim 59, wherein the step of generating said analog signal comprises generating a telephone signal.

61. A method of communicating as recited in claim 60, wherein the step of generating said telephone signal comprises generating a voice signal.

62. A method of communicating as recited in claim 60, wherein the step of generating said telephone signal comprises generating a telephonic fax signal.

63. A method of communicating as recited in claim 60, wherein the step of generating said telephone signal comprises generating a telephonic data modem signal.

64. A method of communicating as recited in claim 59, wherein the step of generating said analog signal comprises generating a video signal.

65. A method of communicating as recited in claim 59, wherein the step of generating said analog signal comprises generating an audio signal.

66. A method of communicating as recited in claim 44, the step of transmitting said at least one first upstream modulated carrier signal wirelessly comprises transmitting said at least one first upstream modulated carrier signal wirelessly from each one of a plurality of subscriber access interface devices.

67. A method of communicating as recited in claim 66, the step of transmitting said at least one first upstream modulated carrier signal wirelessly from each one of said plurality of subscriber access interface devices further comprising the step of arbitrating for access to said network access interface device from said plurality of subscriber access interface devices.

68. A method of communicating as recited in claim 44, the steps of modulating and demodulating said upstream baseband signal onto at least one first upstream radio frequency carrier using a direct sequence spread spectrum process.

69. A method of communicating as recited in claim 44, the steps of modulating and demodulating said upstream baseband signal onto at least one first upstream radio frequency carrier using a frequency hopping spread spectrum process.

70. A method of communicating as recited in claim 44, the steps of modulating and demodulating said upstream baseband signal onto at least one upstream wireless radio frequency carrier using a vector modulation process.

71. A method of communicating as recited in claim 69, the steps of modulating and demodulating said upstream baseband signal onto at least one upstream wireless radio frequency carrier using a quadrature phase shift keying process.

72. A method of communicating as recited in claim 69, the steps of modulating and demodulating said upstream baseband signal onto at least one upstream wireless radio frequency carrier using a bi–phase shift keying process.

73. A method of communicating as recited in claim 44, the steps of modulating and demodulating said upstream baseband signal onto at least one upstream wireless radio frequency carrier using an IEEE 802.11 compliant process.

74. A method of communicating as recited in claim 44, the steps of modulating and demodulating said upstream baseband signal onto at least one upstream wireless radio frequency carrier using a HiperLAN2 compliant process.

75. A method of communicating as recited in claim 44 and further comprising the steps of:

generating a polling signal at said network access interface device,

receiving said polling signal at a subscriber access interface device,

responding to said polling signal by generating a response upstream baseband signal,

transmitting said response upstream baseband signal, and receiving said response upstream baseband signal in said network access interface device.

76. A method of communicating as recited in claim 75 and further comprising the step of forwarding said response upstream baseband signal to a headend on said linear broadband network.

77. A method of communicating as recited in claim 75 and further comprising the step of storing said response upstream baseband signal for retrieval of said response upstream baseband signal upon request.

78. A method of communicating as recited in claim 44, the step of generating an upstream baseband signal further comprising the steps of generating a communication signal in each one of a plurality of said client devices, converting each said communication signal into respective ones of said upstream baseband signals, time division multiplexing each said respective ones of said upstream baseband signals in a subscriber access interface device to produce a multiplexed upstream baseband signal and the step of modulating further comprising the step of modulating said multiplexed baseband signal onto said at least one first upstream radio frequency carrier.

79. A method of communicating as recited in claim 78, wherein the step of converting said communication signal further comprises the step of coupling said upstream baseband signal to said subscriber access interface through a wired connection.

80. A method of communicating as recited in claim 78, wherein the step of converting said communication signal further comprises the step of coupling said

upstream baseband signal to said subscriber access interface device through a wireless connection.

81. A method of communicating as recited in claim 78, each said upstream baseband signal comprising a series of upstream data packets, the method further comprising the steps of prioritizing a launch of each said upstream data packet onto said multiplexed upstream baseband signal to control latency of each said upstream data packet.

82. A method of communicating as recited in claim 81 and further comprising the steps of interpreting a header of each upstream data packet, assigning a priority to each upstream data packet, producing an assigned priority for each upstream data packet, and multiplexing each upstream data packet according to said assigned priority.

83. A method of communicating as recited in claim 38 and further comprising the steps of assigning a priority to each upstream data packet based upon a source port configuration, producing an assigned priority for each upstream data packet, and multiplexing each upstream data packet according to said assigned priority.

84. A method of communicating as recited in claim 44 and further comprising the step of:

controlling a timing of the step of receiving said first upstream modulated carrier signal.

85. A method of communicating as recited in claim 44 and further comprising the step of filtering a spectral range of all of said at least one upstream modulated carrier signal.

86. A method of communicating as recited in claim 44 and further comprising the step of establishing a communications link with an alternate network access interface device in the event of a communication link degradation below a predetermined threshold with said network access interface device.

87. A method of communicating bi-directional information as recited in claim 44 the step of transmitting said second downstream baseband signal further

comprising the step of routing said at least one second downstream baseband signal to a plurality of client devices.

88. A method of communicating bi–directional information as recited in claim 44 wherein said step of transmitting said at least one first upstream modulated carrier and said step of transmitting said at least one second downstream modulated carrier uses a full–duplex process.

89. A method of communicating bi–directional information as recited in claim 44 wherein said step of transmitting said at least one first upstream modulated carrier and said step of transmitting said at least one second downstream modulated carrier uses a half–duplex process.

90. A system for communicating between a client device and a linear broadband network having substantially linear and broadband frequency characteristics comprising:

means for generating an upstream baseband signal,

means for modulating said upstream baseband signal to produce at least one first upstream modulated carrier signal,

means for transmitting said at least one first upstream modulated carrier signal wirelessly,

means for receiving said at least one first upstream modulated carrier signal at a network access interface device coupled to said linear broadband network,

means for demodulating said at least one first upstream modulated carrier signal to produce an upstream demodulated baseband signal, and

means for modulating said upstream demodulated baseband signal onto at least one second upstream radio frequency carrier to produce at least one second upstream modulated carrier signal having a signal format compatible with the linear broadband network.

91. A system for upstream communication comprising:

a linear broadband network having substantially linear and broadband frequency characteristics,

a first upstream modulator that modulates at least one upstream baseband signal received from a client device, said at least one upstream baseband signal being modulated onto at least one upstream wireless radio frequency carrier to produce a first upstream modulated carrier signal,

an upstream transmitter that wirelessly transmits said at least one first upstream modulated carrier signal,

an upstream receiver that receives said at least one first upstream modulated carrier signal,

an upstream demodulator that demodulates said at least one first upstream modulated carrier signal to generate at least one upstream demodulated baseband signal,

a second upstream modulator that modulates said at least one demodulated upstream baseband signal onto at least one upstream linear broadband network radio frequency carrier for transmission onto said linear broadband network.

92. A system for communicating as recited in claim 91 and further comprising a plurality of client devices, each said plurality of client devices generating a one of said at least one upstream baseband signals.

93. A system for communicating as recited in claim 92, and further comprising a multiplexer that multiplexes said plurality of said at least one upstream baseband signals to produce a multiplexed upstream baseband signal.

94. A system for communicating as recited in claim 93, said upstream baseband signal comprising a plurality of upstream data packets, the system further comprising a prioritization device that assigns a priority to each upstream data packet, said priority informing said multiplexer of an appropriate order in which said multiplexer selects each upstream data packet.

95. A system for communicating as recited in claim 91, wherein said client device is a telephone.

96. A system for communicating as recited in claim 91, wherein said client device is a video device.

97. A system for communicating as recited in claim 91, wherein said client device is a computer.

98. A system for communicating as recited in claim 91, wherein said client device is an audio device.

99. A system for communicating as recited in claim 91 and further comprising an arbitrator that controls an order in which said upstream transmitter transmits said at least one first upstream modulated carrier signal.

100. A system for communicating as recited in claim 91 and further comprising an arbitrator that controls an order in which said receiver receives said at least one first upstream modulated carrier signal.

101. A system for communicating as recited in claim 91, and further comprising a forward error correction encoder and a forward error correction decoder.

102. A system for communicating as recited in claim 91 wherein said upstream baseband signal is transmitted over a wired connection.

103. A system for communicating as recited in claim 91 wherein said upstream baseband signal is modulated onto a local upstream carrier and transmitted wirelessly.

104. A system for communicating as recited in claim 91, wherein said first upstream modulator and said upstream transmitter are integral with said client device.

105. A system for communicating as recited in claim 91, wherein said first upstream modulator and said upstream transmitter comprise a peripheral device to said client device.

106. A system for bi–directional communication comprising:

a bi–directional linear broadband network having substantially linear and broadband frequency characteristics,

a first upstream modulator that modulates at least one upstream baseband signal received from a client device, said at least one upstream baseband signal being modulated onto at least one upstream wireless radio frequency carrier to produce at least one first upstream modulated carrier signal,

an upstream transmitter that wirelessly transmits said at least one first upstream modulated carrier signal,

an upstream receiver that receives said at least one first upstream modulated carrier signal,

an upstream demodulator that demodulates said at least one first upstream modulated carrier signal to generate at least one upstream demodulated baseband signal,

a second upstream modulator that modulates said at least one demodulated upstream baseband signal onto at least one upstream linear broadband network radio frequency carrier for transmission onto said linear broadband network,

a first downstream receiver that receives at least one first downstream modulated carrier signal from said linear broadband network,

a first downstream demodulator that demodulates said at least one first downstream modulated carrier signal to produce a first downstream baseband signal,

a downstream modulator that modulates said first downstream baseband signal onto a downstream wireless radio frequency carrier to produce a second downstream modulated carrier signal,

a downstream transmitter that transmits said downstream modulated carrier signal,

a second downstream receiver that receives said downstream modulated carrier signal,

a second downstream demodulator that demodulates said downstream modulated carrier signal to produce a second downstream baseband signal for delivery to said client device.

107. A system for communicating as recited in claim 106 and further comprising a plurality of client devices, each said plurality of client devices generating one of said at least one upstream baseband signals.

108. A system for communicating as recited in claim 107, and further comprising a multiplexer that multiplexes said plurality of said at least one upstream baseband signals to produce a multiplexed upstream baseband signal.

109. A system for communicating as recited in claim 106, wherein said client device is a telephone.

110. A system for communicating as recited in claim 106, wherein said client device is a video device.

111. A system for communicating as recited in claim 106, wherein said client device is a computer.

112. A system for communicating as recited in claim 106, wherein said client device is an audio device.

113. A system for communicating as recited in claim 106 and further comprising an arbitrator that controls an order in which said upstream transmitter transmits said at least one first upstream modulated carrier signal.

114. A system for communicating as recited in claim 106 and further comprising an arbitrator that controls an order in which said upstream receiver receives said at least one upstream modulate carrier signal.

115. A system for communicating as recited in claim 106, and further comprising a forward error correction encoder and a forward error correction decoder.

116. A system for communicating as recited in claim 108, said upstream baseband signal comprising a plurality of upstream data packets, the system further comprising a prioritization device that assigns a priority to each upstream data packet, said priority informing said multiplexer of an appropriate order in which said multiplexer selects each upstream data packet.

117. A system for communicating as recited in claim 106 wherein said upstream baseband signal is transmitted over a wired connection.

118. A system for communicating as recited in claim 106 wherein said upstream baseband signal is modulated onto a local upstream carrier and transmitted wirelessly.

119. A system for communicating as recited in claim 106, wherein said first upstream modulator and said upstream transmitter are integral with said client device.

120. A system for communicating as recited in claim 106, wherein said first upstream modulator and said upstream transmitter comprise a peripheral device to said client device.

121. An apparatus for coupling to a linear broadband network, the linear broadband network having substantially linear and broadband frequency characteristics comprising:

an upstream receiver that wirelessly receives at least one first upstream modulated carrier signal,

an upstream demodulator that demodulates said at least one upstream modulated carrier signal to produce at least one demodulated upstream baseband signal,

an upstream modulator that modulates said at least one demodulated upstream baseband signal onto at least one upstream linear broadband network radio frequency carrier to produce at least one second upstream modulated carrier signal for transmission on the linear broadband network, and

an upstream transmitter that transmits said at least one second upstream modulated carrier signal onto the linear broadband network.

122. An apparatus for coupling as recited in claim 121 and further comprising an arbitrator for controlling a timing of receipt of said at least one upstream radio frequency carrier from a plurality of client devices.

123. An apparatus for coupling as recited in claim 121 and further comprising a polling device for acquiring information from a client device.

124. An apparatus for communicating with a linear broadband network having substantially linear and broadband frequency characteristics comprising:

an upstream receiver for receiving a plurality of upstream baseband signals over a wired connection from a plurality of client devices,

a multiplexer for multiplexing said plurality of upstream baseband signals onto a multiplexed upstream baseband signal,

a first upstream modulator for modulating said at least one multiplexed upstream baseband signal onto at least one upstream wireless radio frequency carrier, and

an upstream transmitter for transmitting said at least one upstream wireless radio frequency wireless carrier.

125. An apparatus for communicating as recited in claim 124 wherein said at least one upstream baseband signal comprises a plurality of upstream data packets, the apparatus further comprising a prioritization device that assigns a priority to each one of said upstream data packets, said priority informing said multiplexer of an appropriate order in which said multiplexer selects each upstream data packet to produce said multiplexed upstream baseband signal.

126. A system for upstream communication comprising:

a linear broadband network having substantially linear and broadband frequency characteristics,

a subscriber access interface device that receives an upstream baseband signal from a client device, modulates said upstream baseband signal onto at least one upstream wireless radio frequency carrier to produce at least one first upstream modulated carrier signal , and wirelessly transmits said at least one first upstream wireless modulated carrier signal,

a network access interface device, coupled to said linear broadband network, that receives said at least one first upstream modulated carrier signal, demodulates said at least one first upstream modulated carrier signal to produce at least one demodulated upstream baseband signal, modulates said at least one demodulated upstream baseband signal onto at least one upstream linear broadband network radio frequency carrier having a format compatible with said linear broadband network to produce at least one second upstream modulated carrier signal, and transmits said at least one second upstream modulated carrier signal onto said linear broadband network.

127. A system for upstream communication as recited in claim 126 and further comprising a first client device coupled to said subscriber access interface device, wherein said first client device interfaces to a local area network that is connected to a plurality of second client devices.

128. A system for upstream communication as recited in claim 126 and further comprising an alternate network access interface device wherein said subscriber access interface device communicates with said network access interface device until a communication link with said network access interface device falls below a predetermined threshold and upon falling below said predetermined threshold, said subscriber access interface device establishes an alternate communication link with said alternate network access interface device.

129. A system for upstream communication as recited in claim 126 and further comprising an alternate network access interface device wherein said subscriber access interface device communicates with said network access interface device and said alternate network access interface device according to the existence of an adequate communication link.

130. A system for upstream communication comprising:

a linear broadband network having substantially linear and broadband frequency characteristics,

a subscriber access interface device that receives an upstream baseband signal, modulates said upstream baseband signal, and wirelessly transmits said modulated upstream baseband signal,

a network access interface device, coupled to said linear broadband network, that receives said modulated upstream baseband signal, demodulates said modulated upstream baseband signal to produce at least one demodulated upstream baseband signal, modulates said at least one demodulated upstream baseband signal onto at least one upstream linear broadband network radio frequency carrier having a format compatible with said linear broadband network to produce at least one second upstream modulated carrier signal, and transmits said at least one second upstream modulated carrier signal onto said linear broadband network.

131. A method of communicating information from a client device to a linear broadband network having substantially linear and broadband frequency characteristics comprising the steps of:

generating an upstream baseband signal,

modulating the upstream baseband signal onto an upstream wireless radio frequency carrier to generate a first upstream modulated carrier signal,

transmitting said first upstream modulated carrier signal wirelessly,

receiving said first upstream modulated carrier signal at a network access interface device coupled to the linear broadband network,

demodulating said first upstream modulated carrier signal to produce an upstream demodulated baseband signal, and

modulating said upstream demodulated baseband signal onto an upstream linear broadband radio frequency carrier to produce a second upstream modulated carrier signal that is compatible with the linear broadband network.

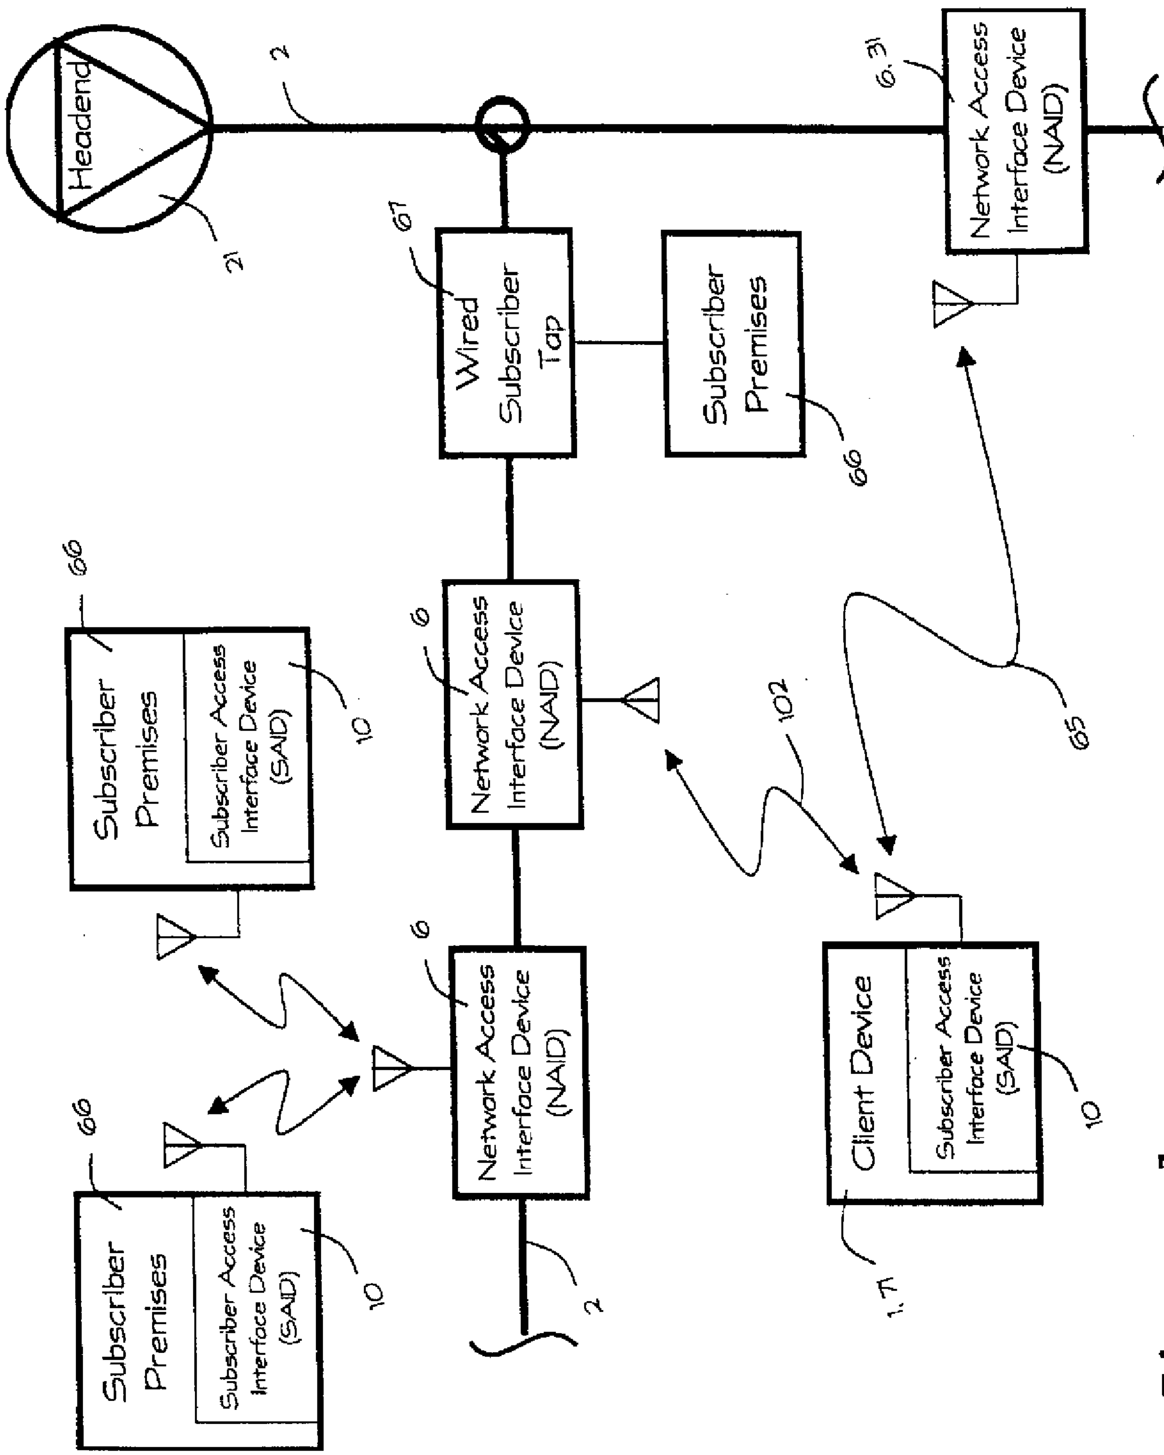
Headend
21
2
6, 31
Network Access
Interface Device
(NAID)
67
Wired
Subscriber
Tap
Subscriber
Premises
66
Subscriber
Premises
Subscriber Access
Interface Device
(SAID)
10
6
Network Access
Interface Device
(NAID)
102
65
Client Device
1, 71

Figure 1

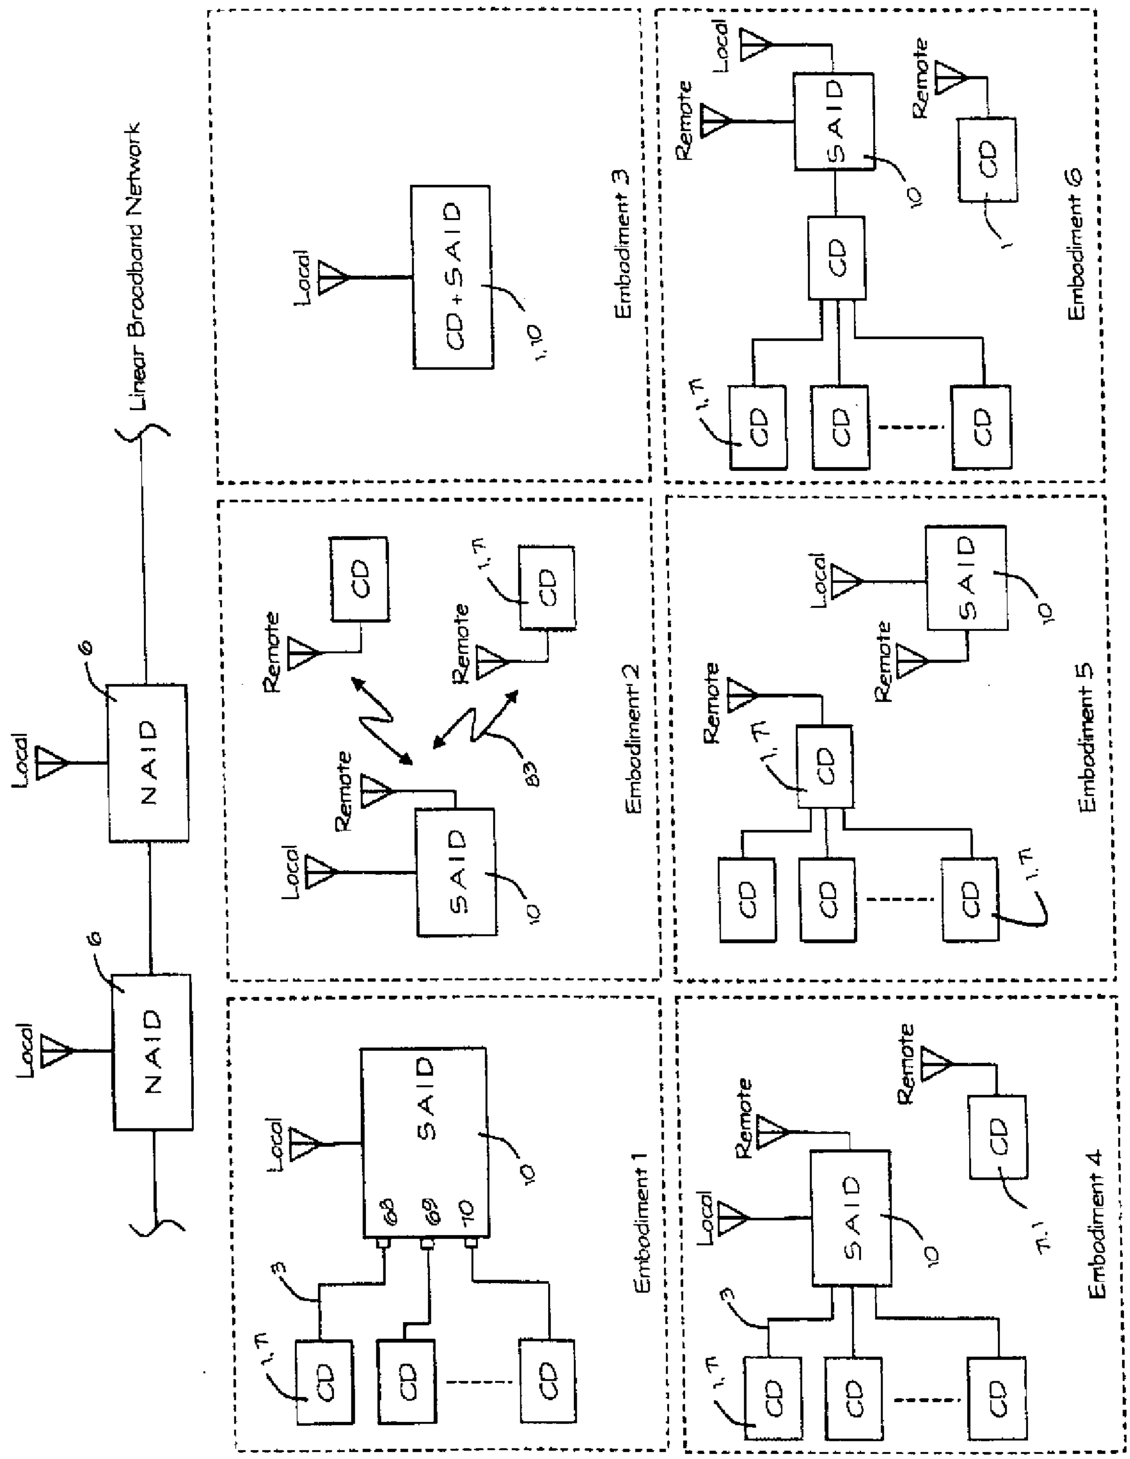
Linear Broadband Network
Local
NAID
6
SAID
CD
CD + SAID
Remote
68
69
70
3
1, 71
1, 10
10
83
71.1
Embodiment 1
Embodiment 2
Embodiment 3
Embodiment 4
Embodiment 5
Embodiment 6

Figure 2

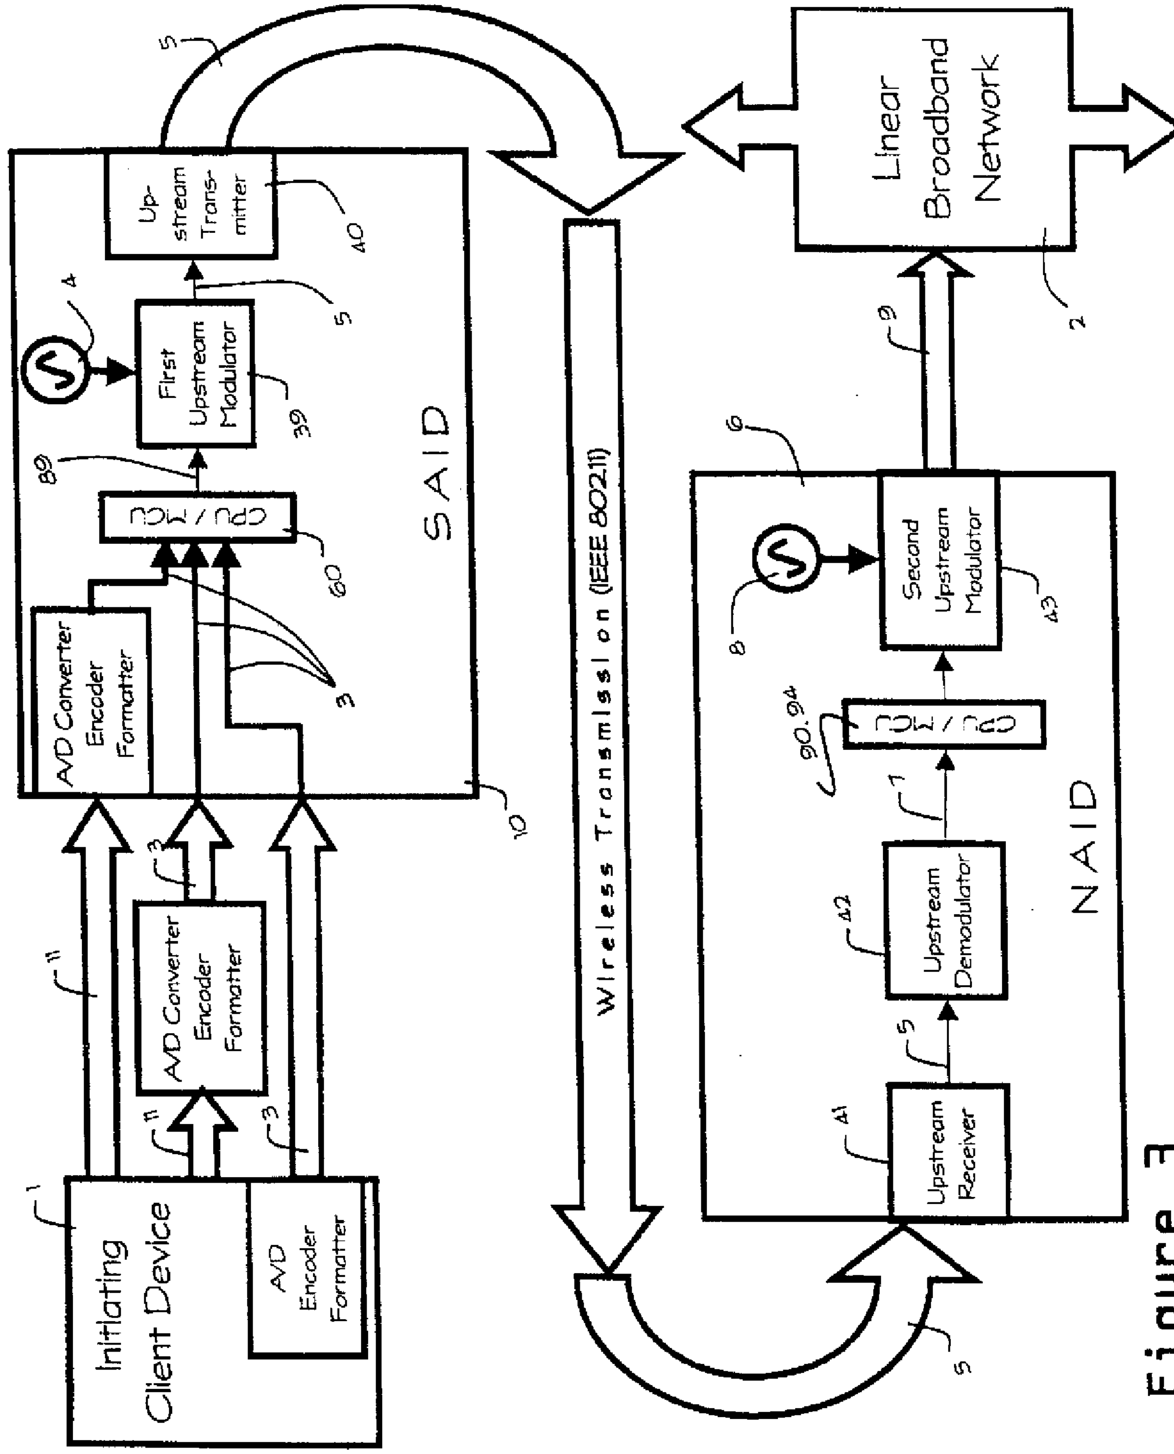
Initiating Client Device
A/D Encoder Formatter
A/D Converter Encoder Formatter
A/D Converter Encoder Formatter
CPU / MCU
First Upstream Modulator
Up-stream Trans-mitter
SAID
Wireless Transmission (IEEE 802.11)
Linear Broadband Network
Second Upstream Modulator
CPU / MCU
Upstream Demodulator
Upstream Receiver
NAID
1
11
11
3
3
3
10
60
89
4
39
5
40
5
5
6
8
43
9
2
90,94
7
42
5
41

Figure 3

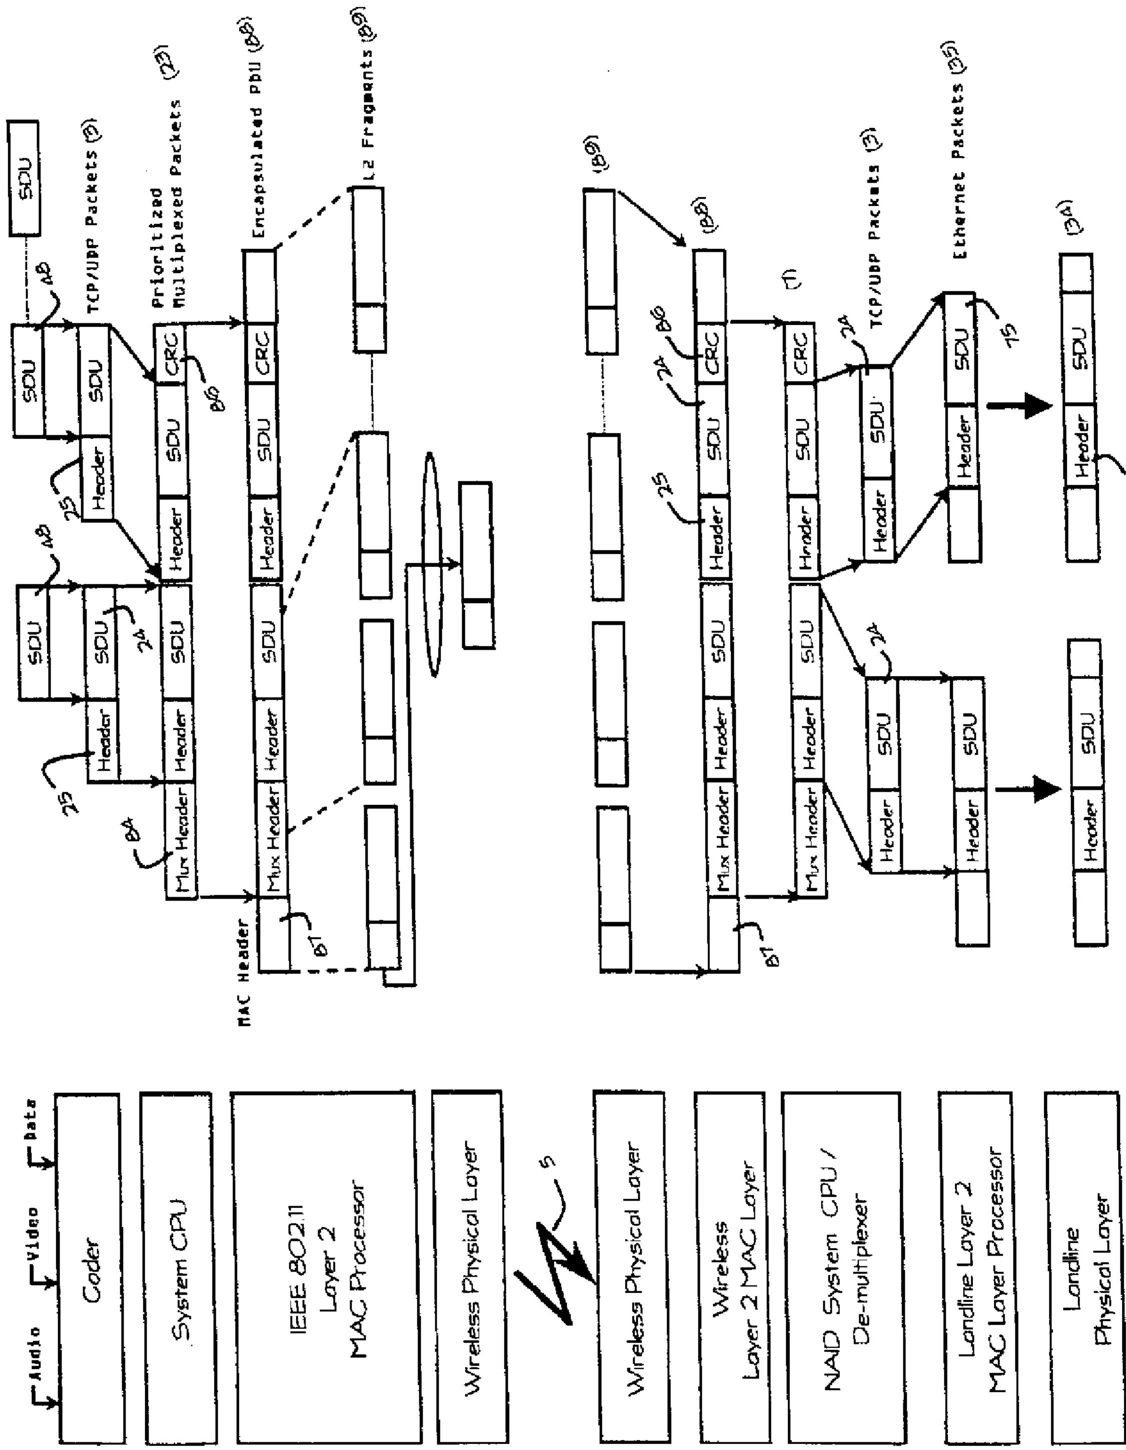
Audio
Video
Data
Coder
System CPU
IEEE 802.11 Layer 2 MAC Processor
Wireless Physical Layer
5
Wireless Physical Layer
Wireless Layer 2 MAC Layer
NAID System CPU / De-multiplexer
Landline Layer 2 MAC Layer Processor
Landline Physical Layer
SDU
TCP/UDP Packets (3)
Prioritized Multiplexed Packets (23)
Encapsulated PDU (88)
L2 Fragments (89)
MAC Header
Mux Header
Header
CRC
Ethernet Packets (35)
(34)

Figure 4

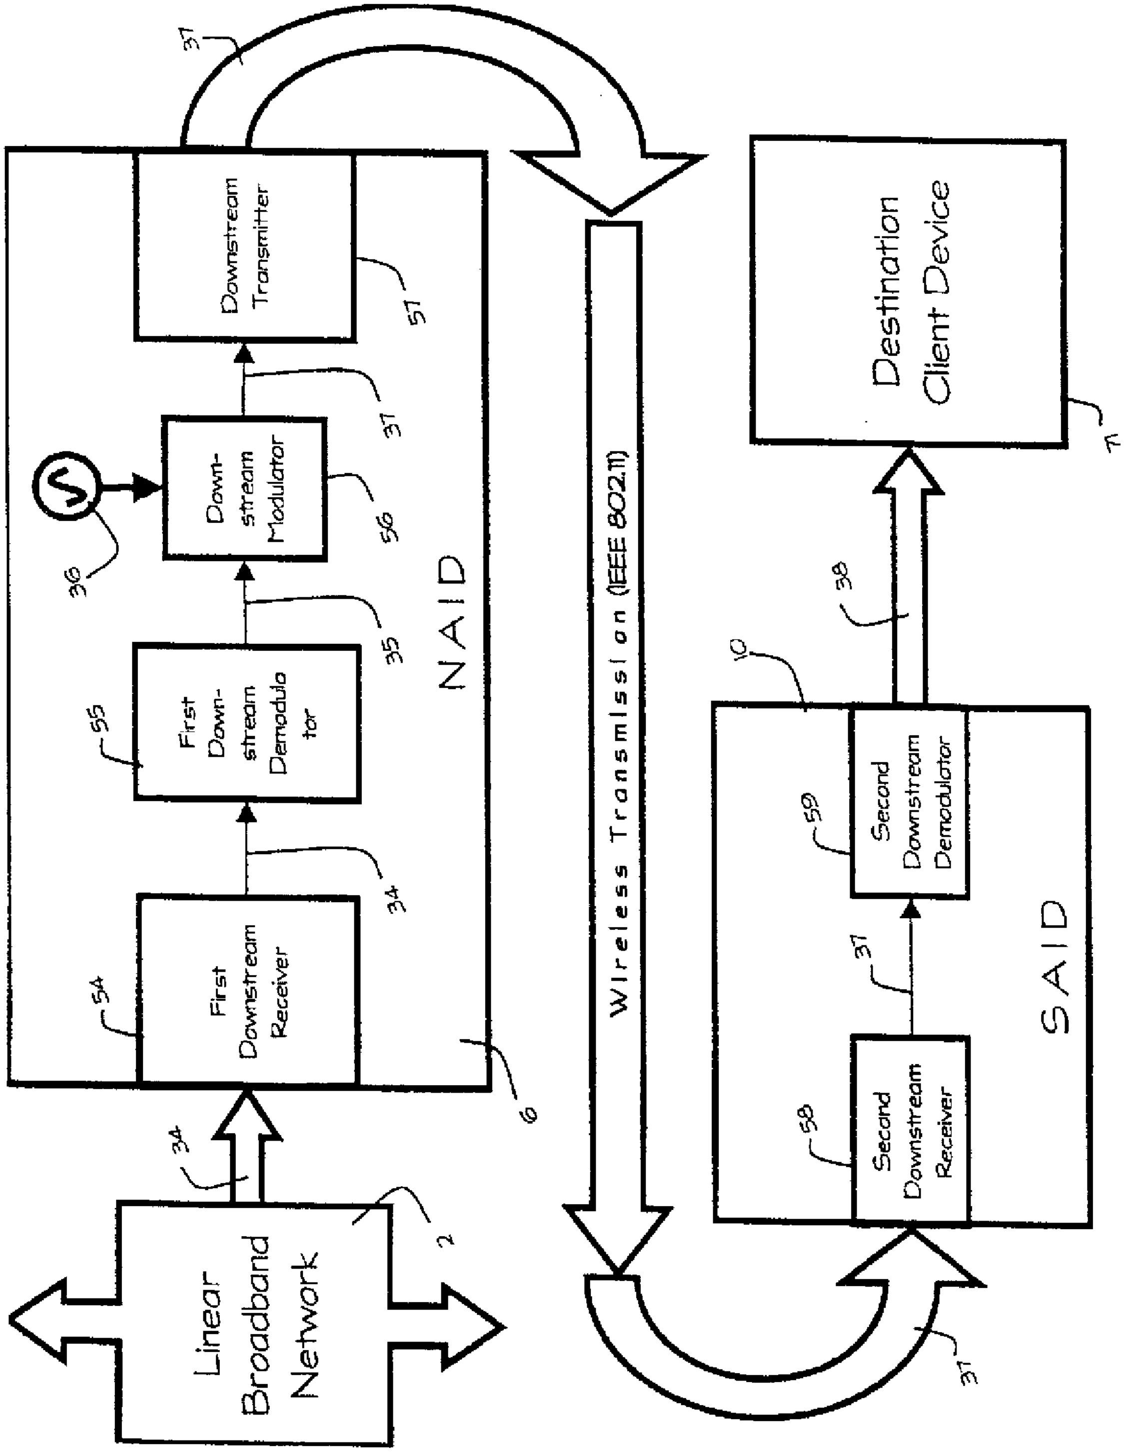
Linear Broadband Network
2
34
First Downstream Receiver
54
34
First Down-stream Demodula tor
55
35
Down-stream Modulator
56
36
37
Downstream Transmitter
57
37
NAID
6
Wireless Transmission (IEEE 802.11)
37
Second Downstream Receiver
58
37
Second Downstream Demodulator
59
SAID
10
38
Destination Client Device
7

Figure 5

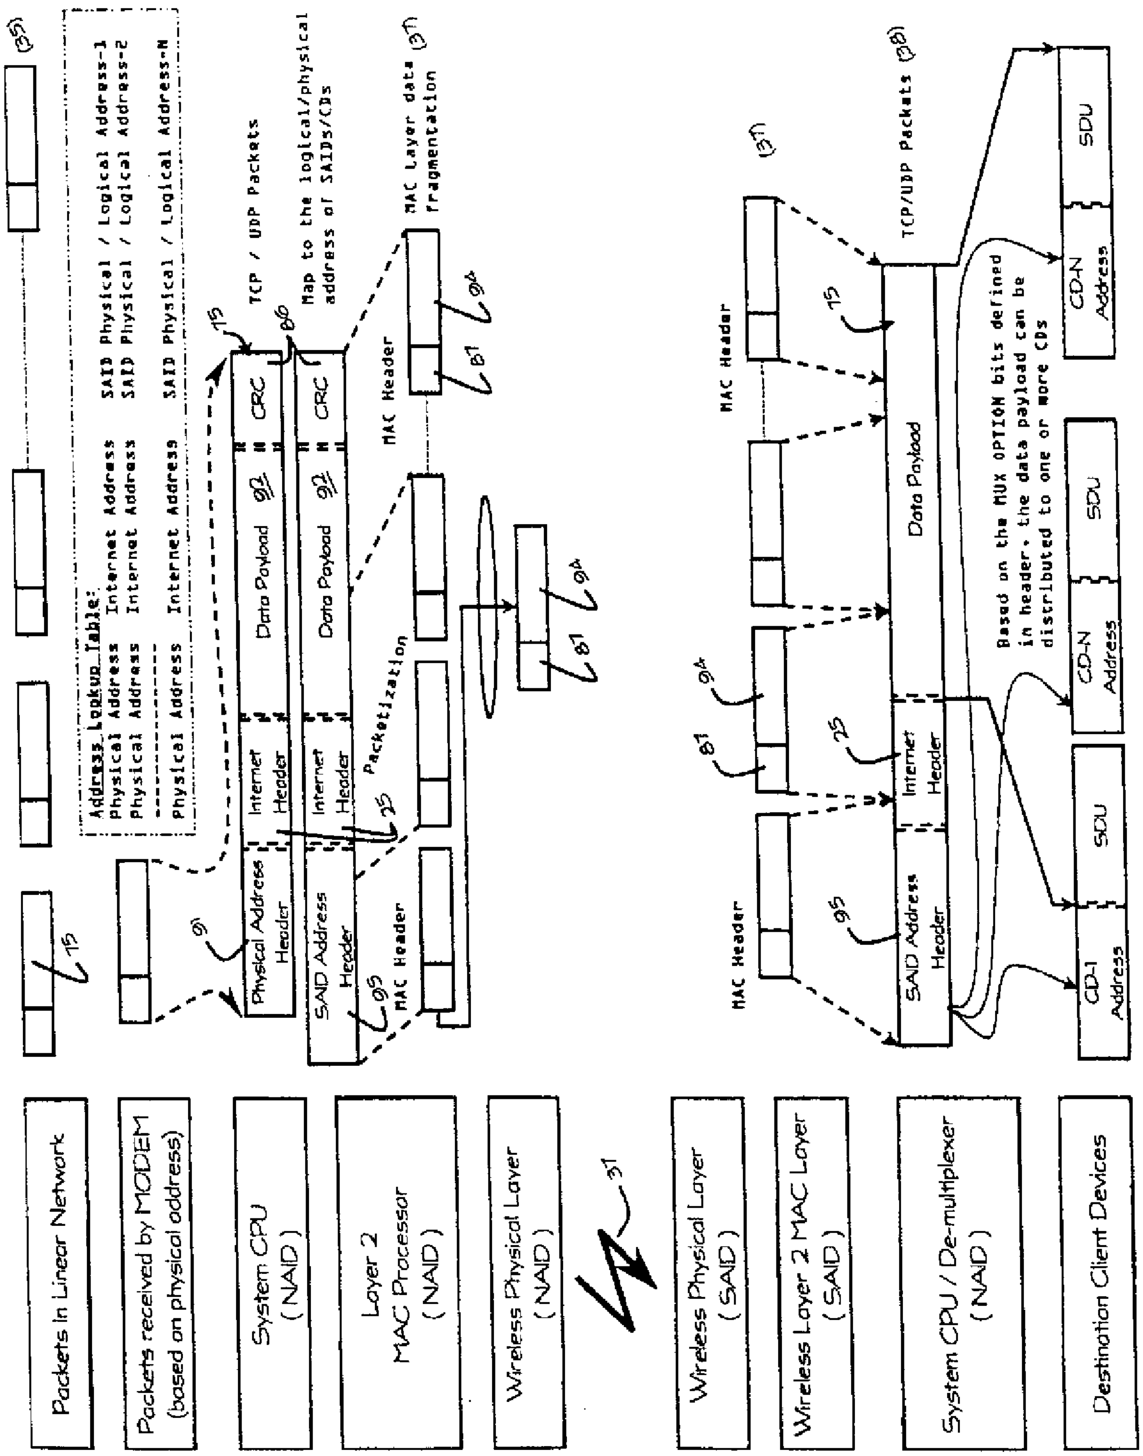
Packets In Linear Network
Packets received by MODEM (based on physical address)
System CPU (NAID)
Layer 2 MAC Processor (NAID)
Wireless Physical Layer (NAID)
31
Wireless Physical Layer (SAID)
Wireless Layer 2 MAC Layer (SAID)
System CPU / De-multiplexer (NAID)
Destination Client Devices
(35)
15
Address Lookup Table:
Physical Address Internet Address SAID Physical / Logical Address-1
Physical Address Internet Address SAID Physical / Logical Address-2
Physical Address Internet Address SAID Physical / Logical Address-N
91
Physical Address Header
Internet Header
Data Payload 92
CRC
86
TCP / UDP Packets
SAID Address Header
95
25
Map to the logical/physical address of SAIDs/CDs
MAC Header
Packetization
MAC Layer data fragmentation
(37)
81
94
Data Payload
TCP/UDP Packets (38)
Based on the MUX OPTION bits defined in header, the data payload can be distributed to one or more CDs
CD-1 Address
SDU
CD-N Address

Figure 6

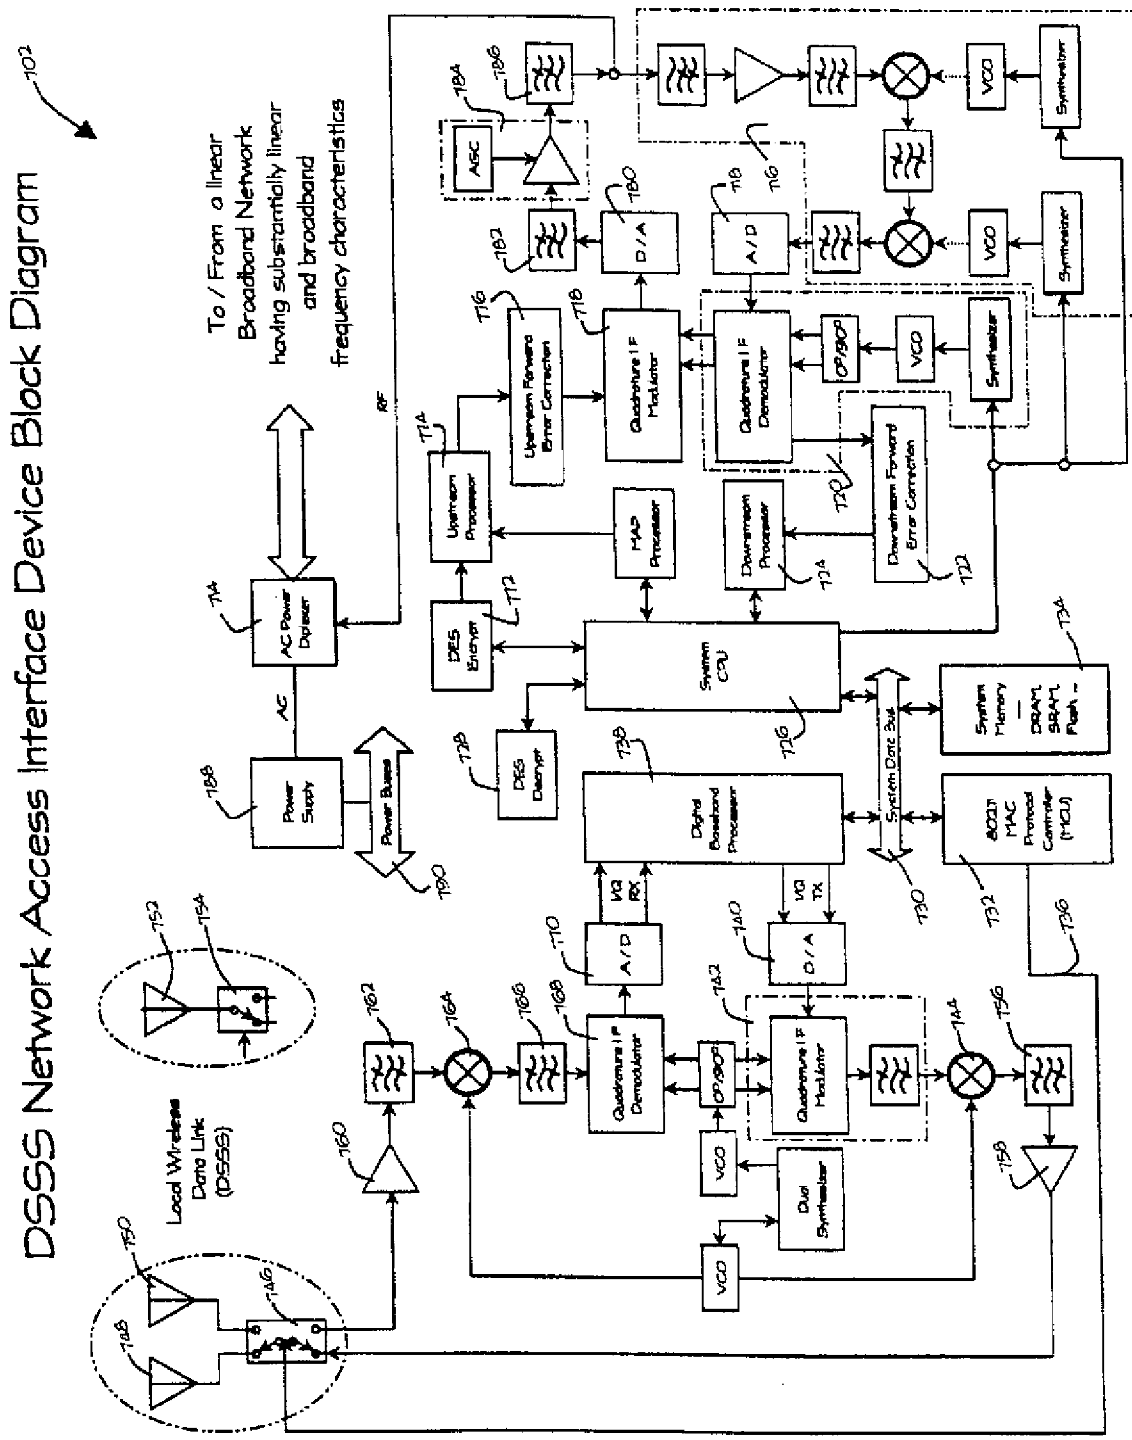
DSSS Network Access Interface Device Block Diagram
102
To / From a linear Broadband Network having substantially linear and broadband frequency characteristics
Local Wireless Data Link (DSSS)
AC Power Diplexer
Power Supply
Power Buses
AC
RF
DES Encryptor
DES Decryptor
System CPU
Upstream Processor
MAP Processor
Downstream Processor
Upstream Forward Error Correction
Downstream Forward Error Correction
Quadrature IF Modulator
Quadrature IF Demodulator
D / A
A / D
AGC
VCO
Synthesizer
Dual Synthesizer
Digital Baseband Processor
System Data Bus
System Memory DRAM, SRAM, Flash
802.11 MAC Protocol Controller (MCU)
I/Q RX
I/Q TX

Figure 7

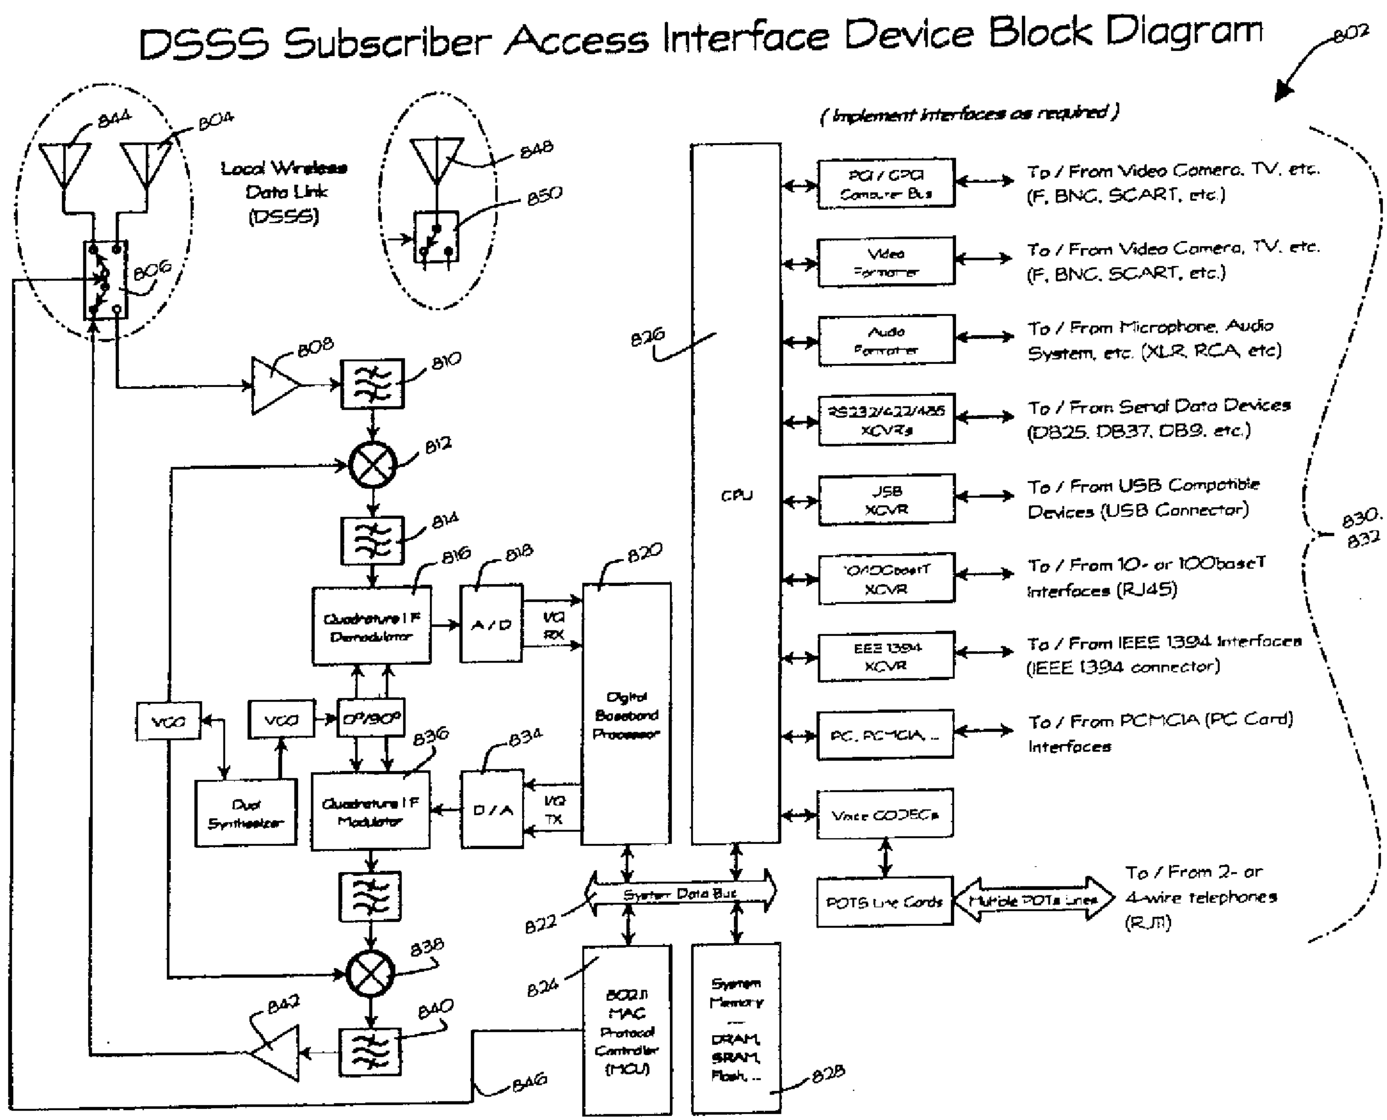
DSSS Subscriber Access Interface Device Block Diagram
802
(Implement interfaces as required)
844
804
Local Wireless
Data Link
(DSSS)
848
850
806
808
810
812
814
816
818
820
826
CPU
PCI / CPCI
Computer Bus
Video
Formatter
Audio
Formatter
RS232/422/485
XCVRs
USB
XCVR
10/100baseT
XCVR
IEEE 1394
XCVR
PC, PCMCIA, ...
Voice CODECs
POTS Line Cards
Multiple POTs Lines
To / From Video Camera, TV, etc.
(F, BNC, SCART, etc.)
To / From Video Camera, TV, etc.
(F, BNC, SCART, etc.)
To / From Microphone, Audio
System, etc. (XLR, RCA, etc)
To / From Serial Data Devices
(DB25, DB37, DB9, etc.)
To / From USB Compatible
Devices (USB Connector)
To / From 10- or 100baseT
Interfaces (RJ45)
To / From IEEE 1394 Interfaces
(IEEE 1394 connector)
To / From PCMCIA (PC Card)
Interfaces
To / From 2- or
4-wire telephones
(RJ11)
830,
832
Quadrature I F
Demodulator
A / D
I/Q
RX
Digital
Baseband
Processor
VCO
VCO
0°/90°
836
834
Dual
Synthesizer
Quadrature I F
Modulator
D / A
I/Q
TX
System Data Bus
822
838
824
842
840
802.11
MAC
Protocol
Controller
(MCU)
System
Memory
DRAM,
SRAM,
Flash, ...
828
846

Figure 8

FHSS Network Access Interface Device Block Diagram

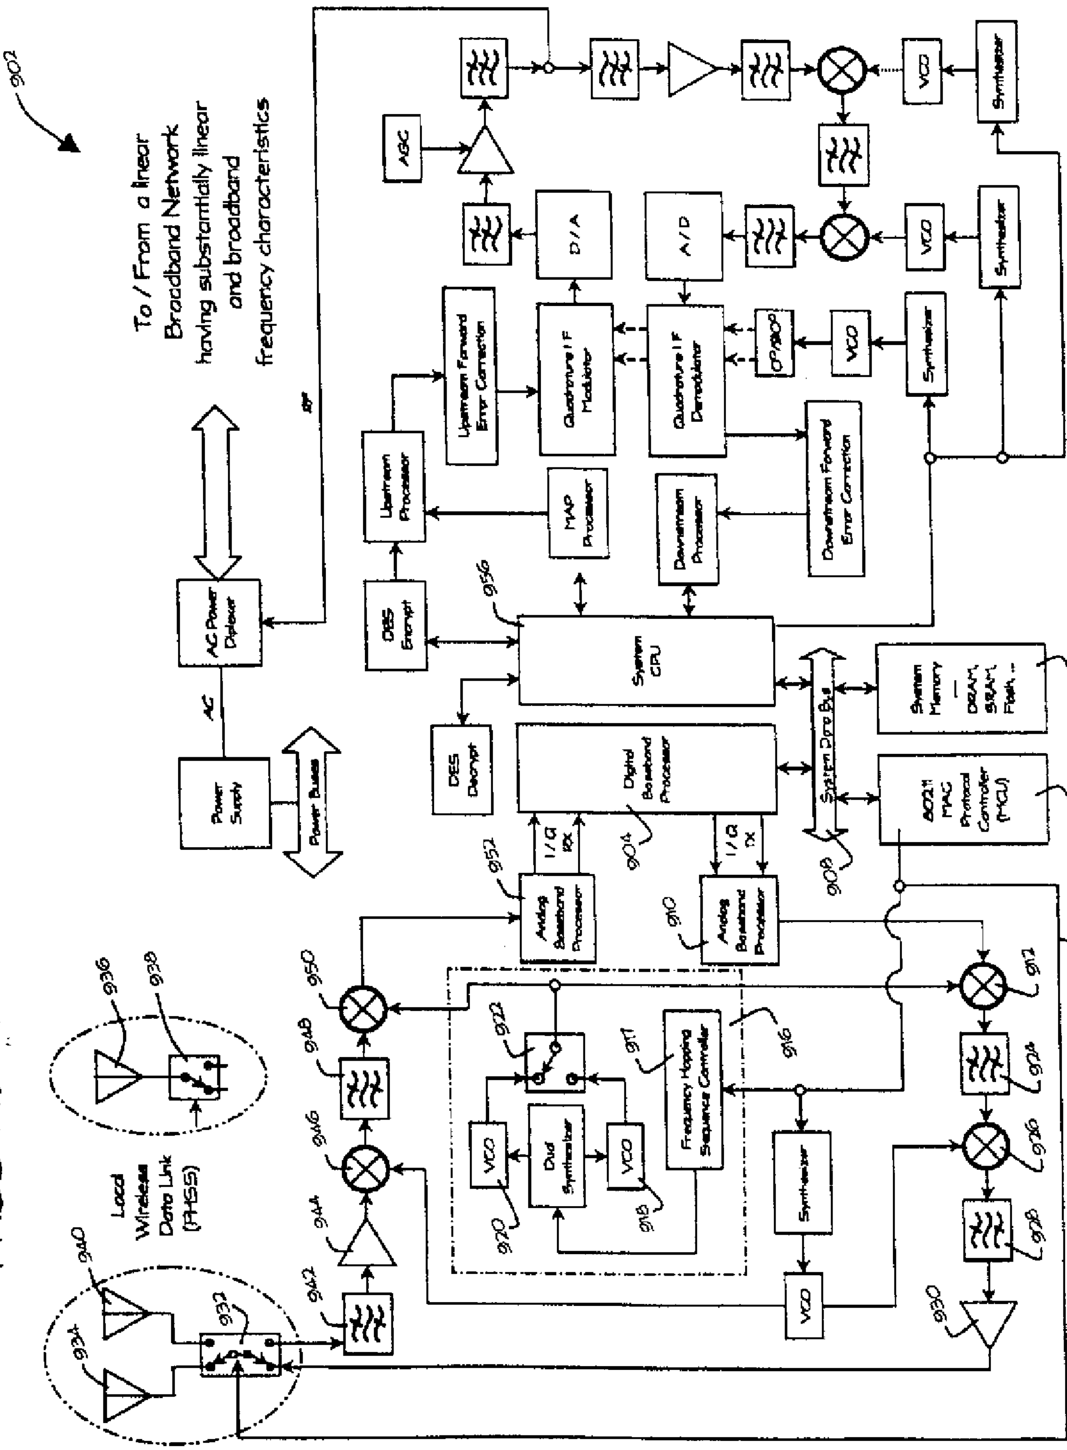

Figure 9

FHSS Subscriber Access Interface Device Block Diagram

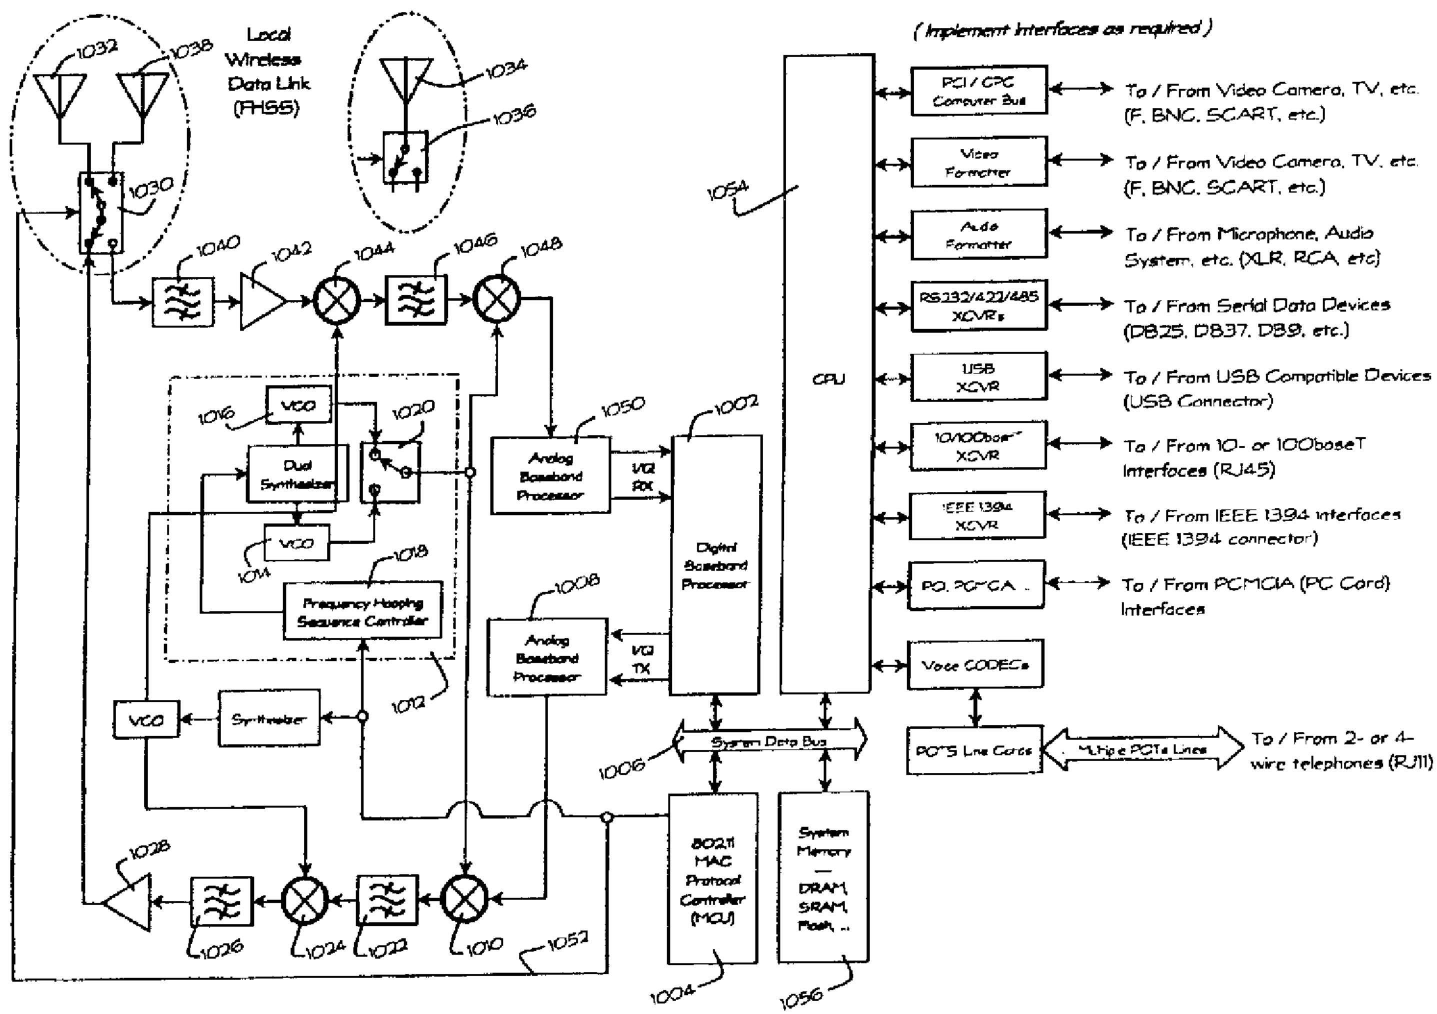

Figure 10

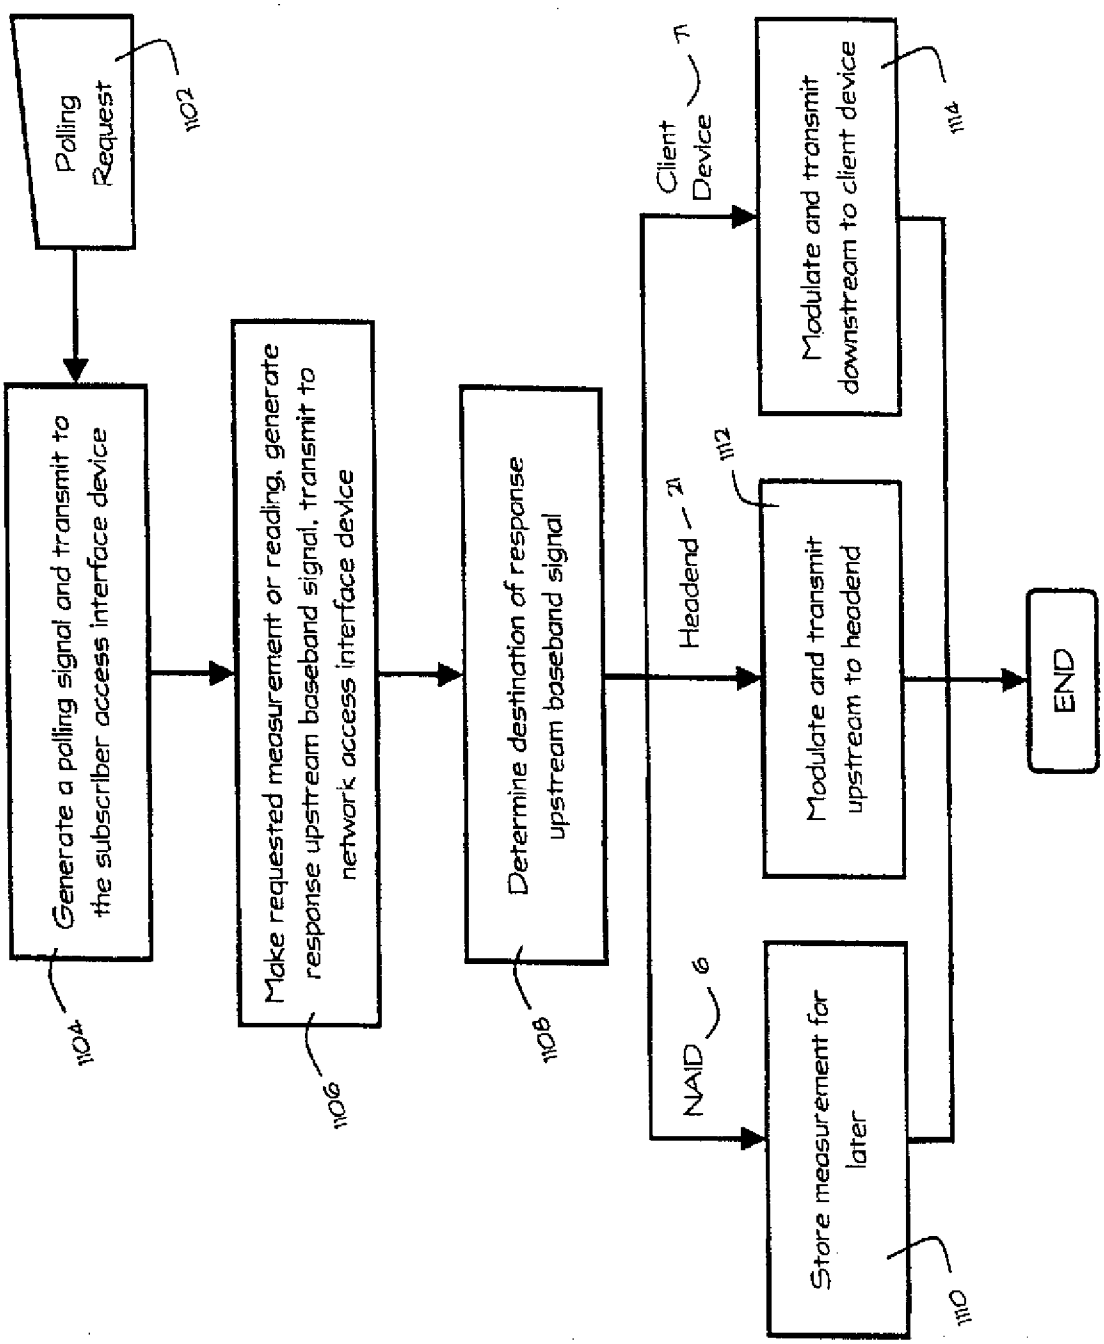
Polling Request
1102
Generate a polling signal and transmit to the subscriber access interface device
1104
Make requested measurement or reading, generate response upstream baseband signal, transmit to network access interface device
1106
Determine destination of response upstream baseband signal
1108
NAID
6
Headend
21
Client Device
71
Store measurement for later
1110
Modulate and transmit upstream to headend
1112
Modulate and transmit downstream to client device
1114
END

Figure 11

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
8 September 2000 (08.09.2000)

PCT

(10) International Publication Number
WO 00/52880 A3

(51) International Patent Classification[7]: H04L 12/28, 1/00, 12/403, 12/56

(21) International Application Number: PCT/US00/05270

(22) International Filing Date: 29 February 2000 (29.02.2000)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:

60/122,287	1 March 1999 (01.03.1999)	US
60/128,268	8 April 1999 (08.04.1999)	US
60/142,082	29 June 1999 (29.06.1999)	US
09/505,107	15 February 2000 (15.02.2000)	US

(71) Applicant *(for CA only)*: **MEDIACELL, INC.** [US/US]; 16280 NW Mission Oaks Drive, Beaverton, OR 97006 (US).

(71) Applicant *(for all designated States except CA)*: **PACIFIC COMMUNICATIONS LIMITED** [CN/CN]; 8/F Bank of America Tower, 12 Harcourt Road, Central, Hong Kong (CN).

(72) Inventors: **BISHOP, Donald, M.**; 16280 N.W. Mission Oaks Drive, Beaverton, OR 97006 (US). **GREEN, George, R., J.**; 39078 Guardino Drive #102, Fremont, CA 94538 (US). **SHYU, Archie, R.**; 12460 Dormouse, San Diego, CA 92129 (US).

(74) Agent: **BRUESS, Steven, C.**; Merchant & Gould P.C., 3100 Norwest Center, 90 South Seventh Street, Minneapolis, MN 55402-4131 (US).

(81) Designated States *(national)*: AE, AL, AM, AT, AT (utility model), AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CR, CU, CZ, CZ (utility model), DE, DE (utility model), DK, DK (utility model), DM, EE, EE (utility model), ES, FI, FI (utility model), GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KR (utility model), KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SK (utility model), SL, TJ, TM, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZW.

(84) Designated States *(regional)*: ARIPO patent (GH, GM, KE, LS, MW, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU,

[Continued on next page]

(54) Title: METHOD AND APPARATUS FOR COMMUNICATING BETWEEN A CLIENT DEVICE AND A LINEAR BROADBAND NETWORK

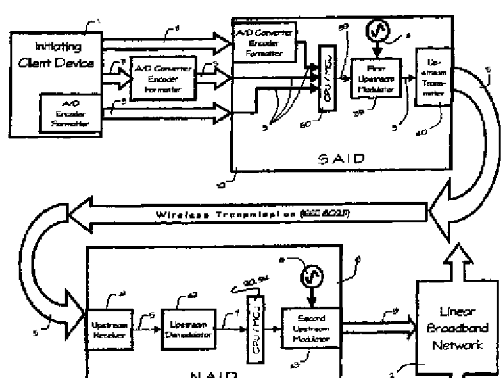

(57) Abstract: A method of upstream communication over a linear broadband network includes the steps of generating an upstream baseband signal and modulating it onto an upstream wireless radio frequency carrier to produce a first upstream modulated carrier signal. The modulated carrier signal is transmitted wirelessly, received, and demodulated to reproduce the information integrity of the upstream baseband signal. The signal is then modulated onto an upstream linear broadband radio frequency carrier for transmission on the linear broadband network. Advantageously, noise that accumulates at the subscriber premises is removed from the upstream signal prior to presentation of the signal to the upstream path of the linear broadband network. A system for communicating over a linear broadband network includes network access interface devices coupled to the linear broadband network. A subscriber access interface device accepts upstream communication signals and modulates and transmits the signal to the network access interface device. The network access interface device (6) receives and demodulates the signal and then modulates it for transmission on the linear broadband network.

WO 00/52880 A3

MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

Published:
— *With international search report.*

(88) Date of publication of the international search report: 25 January 2001

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

INTERNATIONAL SEARCH REPORT

International Application No
PCT/US 00/05270

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 H04L12/28 H04L1/00 H04L12/403 H04L12/56

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 7 H04L H04Q

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

WPI Data, PAJ, EPO-Internal

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 5 421 030 A (BARAN PAUL) 30 May 1995 (1995-05-30)	1,2,9, 11,12, 16-18, 20-23, 44,45, 52,54, 55, 59-61, 63-66, 87, 90-92, 95-98, 102,104, 106,107, 109-112, 117,119, 121,126, 127,130, 131

-/--

[X] Further documents are listed in the continuation of box C.

[X] Patent family members are listed in annex.

° Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
7 August 2000	10. 11. 2000

Name and mailing address of the ISA	Authorized officer
European Patent Office, P.B. 5818 Patentlaan 2 NL – 2280 HV Rijswijk Tel. (+31-70) 340-2040, Tx. 31 651 epo nl, Fax: (+31-70) 340-3016	Vaskimo, K

4

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category °	Citation of document, with indication,where appropriate, of the relevant passages	Relevant to claim No.
	column 1, line 8 -column 4, line 63 column 5, line 62 -column 6, line 22 column 6, line 66 -column 8, line 48 column 11, line 13 -column 12, line 23 column 18, line 44 - line 54 figures 6,10,11	
Y		3,4,8, 10, 13-15, 19, 25-28, 30,46, 47,51, 53, 56-58, 62, 68-71, 73,101, 103,105, 115,118, 120
A		29,31, 72,74

Y	WO 94 15294 A (SUREFIND CORP ;LALONDE JAMES E (US); RUFF RITCHEY A (US)) 7 July 1994 (1994-07-07) abstract page 1, line 8 -page 9, line 9	3,46
A		1,44

Y	US 5 463 656 A (POLIVKA ALAN L ET AL) 31 October 1995 (1995-10-31) column 1, line 10 -column 4, line 18 column 6, line 10 - line 65 column 7, line 27 -column 8, line 12 figures 2A,2B	4,47, 101,115
A		1,44,91, 106

Y	US 5 790 806 A (KOPERDA FRANCIS RICHARD) 4 August 1998 (1998-08-04) abstract column 1, line 4 -column 3, line 3 column 4, line 45 -column 5, line 23 column 8, line 1 -column 9, line 47 claims 1,12 figures 2,4	4,47, 101,115
A		1,44,91, 106

	-/--	

4

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication,where appropriate, of the relevant passages	Relevant to claim No.
Y	PATENT ABSTRACTS OF JAPAN vol. 1997, no. 03, 31 March 1997 (1997-03-31) & JP 08 298688 A (TOSHIBA CORP), 12 November 1996 (1996-11-12) abstract -& US 5 951 709 A (TANAKA HIROKAZU) 14 September 1999 (1999-09-14) column 1, line 6 -column 5, line 47 column 7, line 4 - line 64 figure 2	4,47, 101,115
A		1,44,91, 106

Y	US 5 610 916 A (SISTANIZADEH KAMRAN ET AL) 11 March 1997 (1997-03-11) column 1, line 15 -column 8, line 61 column 13, line 17 -column 14, line 26 column 22, line 62 -column 24, line 39 column 25, line 66 -column 27, line 3	8,10,13, 15,27, 28,51, 53,56, 70,71
A		1,9,11, 14,29, 57,58

Y	US 5 815 794 A (WILLIAMS THOMAS H) 29 September 1998 (1998-09-29) column 1, line 8 -column 4, line 8 column 4, line 29 -column 5, line 24 column 10, line 18 - line 34	8,10,13, 15,27, 28,51, 53,56, 70,71
A		1,9,11, 14,29, 57,58,72

Y	WO 97 48191 A (BUGAS DMITRY ;GRAVES H DAVID (CA); BROADBAND NETWORKS INC (CA); SC) 18 December 1997 (1997-12-18) page 1, line 3 - line 17 page 16, line 4 -page 17, line 15	25,26, 68,69
A		1,44

Y	US 5 272 700 A (HANSEN DAVID E ET AL) 21 December 1993 (1993-12-21) column 1, line 15 -column 2, line 55 column 3, line 23 -column 4, line 40	27,28, 70,71
A		1,44

	-/--	

International Application No
PCT/US 00/05270

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 5 825 776 A (MOON BILLY G) 20 October 1998 (1998-10-20) column 1, line 7 -column 3, line 63 column 5, line 6 - line 23	30,73
A		1,44
Y	US 5 373 503 A (CHEN KWANG-CHENG) 13 December 1994 (1994-12-13) column 1, line 7 -column 4, line 2	30,73
A		1,44
Y	WO 98 39879 A (NORTHERN TELECOM LTD) 11 September 1998 (1998-09-11) page 1, line 1 -page 3, line 9	13,15, 19,56,62
A		1,14, 16-18, 20,57, 59-61,63
Y	WO 98 15143 A (NOKIA MOBILE PHONES LTD ;SORMUNEN TONI (FI); HALMINEN HARRI (FI);) 9 April 1998 (1998-04-09) abstract page 1, line 1 -page 7, line 13 page 8, line 18 - line 29 page 14, line 18 -page 16, line 2 page 17, line 7 -page 19, line 13	13-15, 56-58, 103,105, 118,120
A		1,44,91, 104,106
Y	EP 0 790 749 A (TEXAS INSTRUMENTS INC) 20 August 1997 (1997-08-20) abstract column 1, line 1 -column 2, line 21 column 13, line 49 -column 20, line 11 claims 1,3,8,11,12,15	13,14, 56,57, 103,105, 118,120
A		1,15,44, 91,104, 106,117

4

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US 00/05270

Box I Observations where certain claims were found unsearchable (Continuation of item 1 of first sheet)

This International Search Report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claims Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:

2. ☐ Claims Nos.:
because they relate to parts of the International Application that do not comply with the prescribed requirements to such an extent that no meaningful International Search can be carried out, specifically:

3. ☐ Claims Nos.:
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box II Observations where unity of invention is lacking (Continuation of item 2 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

see additional sheet

1. ☐ As all required additional search fees were timely paid by the applicant, this International Search Report covers all searchable claims.

2. ☐ As all searchable claims could be searched without effort justifying an additional fee, this Authority did not invite payment of any additional fee.

3. ☐ As only some of the required additional search fees were timely paid by the applicant, this International Search Report covers only those claims for which fees were paid, specifically claims Nos.:

4. ☒ No required additional search fees were timely paid by the applicant. Consequently, this International Search Report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

1-4,8-23,25-31.44-47,51-66.68-74,87,90-92,95-98,101-107,109-112,115,117-121,126,127,130,131

Remark on Protest

☐ The additional search fees were accompanied by the applicant's protest.

☐ No protest accompanied the payment of additional search fees.

FURTHER INFORMATION CONTINUED FROM PCT/ISA/ 210

This International Searching Authority found multiple (groups of) inventions in this international application, as follows:

1. Claims: 1-4,8-23,25-31,44-47,51-66,68-74,87,90-92,95-98, 101-107,109-112,115,117-121,126,127,130,131

   A communication method from a client to a broadband network, where an upstream baseband signal is generated, where said upstream signal is modulated onto an upstream wireless radio frequency carrier at a subscriber access interface, where said upstream modulated carrier signal is transmitted wirelessly, where said transmitted upstream modulated carrier signal is received at a network access interface, where said received upstream modulated carrier signal is demodulated to produce an upstream demodulated baseband signal, where said upstream demodulated baseband signal is modulated to produce a second upstream modulated carrier signal compatible with said broadband network, and where forward error correction is performed.

2. Claims: 1,5-7,42,44,48-50,85

   A communication method from a client to a broadband network, where an upstream baseband signal is generated, where said upstream signal is modulated onto an upstream wireless radio frequency carrier at a subscriber access interface, where said upstream modulated carrier signal is transmitted wirelessly, where said transmitted upstream modulated carrier signal is received at a network access interface, where said received upstream modulated carrier signal is demodulated to produce an upstream demodulated baseband signal, where said upstream demodulated baseband signal is modulated to produce a second upstream modulated carrier signal compatible with said broadband network, and where filtering of modulated carrier signal is carried out.

3. Claims: 1,24,32-34,41,44,67,75-77,84,91,99,100,106,113,114, 121-123

   A communication method from a client to a broadband network, where an upstream baseband signal is generated, where said upstream signal is modulated onto an upstream wireless radio frequency carrier at a subscriber access interface, where said upstream modulated carrier signal is transmitted wirelessly, where said transmitted upstream modulated carrier signal is received at a network access interface, where said received upstream modulated carrier signal is demodulated to produce an upstream demodulated baseband signal, where said upstream demodulated baseband signal is modulated to produce a second upstream modulated carrier signal compatible with said broadband network, and where

FURTHER INFORMATION CONTINUED FROM PCT/ISA/ 210

special access control methods are used to access a network access interface.

4. Claims: 1,35-40,44,78-83,91-94,106-108,116,124,125

A communication method from a client to a broadband network, where an upstream baseband signal is generated, where said upstream signal is modulated onto an upstream wireless radio frequency carrier at a subscriber access interface, where said upstream modulated carrier signal is transmitted wirelessly, where said transmitted upstream modulated carrier signal is received at a network access interface, where said received upstream modulated carrier signal is demodulated to produce an upstream demodulated baseband signal, where said upstream demodulated baseband signal is modulated to produce a second upstream modulated carrier signal compatible with said broadband network, and where upstream traffic from multiple clients is time division multiplexed at a subscriber access interface.

5. Claims: 1,43,44,86,126,128,129

A communication method from a client to a broadband network, where an upstream baseband signal is generated, where said upstream signal is modulated onto an upstream wireless radio frequency carrier at a subscriber access interface, where said upstream modulated carrier signal is transmitted wirelessly, where said transmitted upstream modulated carrier signal is received at a network access interface, where said received upstream modulated carrier signal is demodulated to produce an upstream demodulated baseband signal, where said upstream demodulated baseband signal is modulated to produce a second upstream modulated carrier signal compatible with said broadband network, and where an alternate network access interface is arranged.

6. Claims: 44,88,89

A communication method from a client to a broadband network, where an upstream baseband signal is generated, where said upstream signal is modulated onto an upstream wireless radio frequency carrier at a subscriber access interface, where said upstream modulated carrier signal is transmitted wirelessly, where said transmitted upstream modulated carrier signal is received at a network access interface, where said received upstream modulated carrier signal is demodulated to produce an upstream demodulated baseband signal, where said upstream demodulated baseband signal is modulated to produce a second upstream modulated carrier signal compatible with said broadband network, and where upstream and downstream traffic is transmitted as duplex communication.

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No
PCT/US 00/05270

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 5421030 A	30-05-1995	AU 2782092 A	27-04-1993
		EP 0604592 A	06-07-1994
		JP 6510894 T	01-12-1994
		WO 9306669 A	01-04-1993
WO 9415294 A	07-07-1994	US 5283731 A	01-02-1994
		AU 4405093 A	24-10-1994
		AU 6015594 A	19-07-1994
		WO 9423383 A	13-10-1994
		AU 4405393 A	19-07-1994
		WO 9415428 A	07-07-1994
US 5463656 A	31-10-1995	NONE	
US 5790806 A	04-08-1998	AU 2219597 A	22-10-1997
		WO 9737493 A	09-10-1997
JP 08298688 A	12-11-1996	CN 1144442 A	05-03-1997
		US 5951709 A	14-09-1999
US 5610916 A	11-03-1997	US 5651010 A	22-07-1997
		AU 5314996 A	29-11-1996
		AU 5365496 A	02-10-1996
		WO 9628902 A	19-09-1996
		WO 9637059 A	21-11-1996
		US 5784683 A	21-07-1998
		AU 5424496 A	02-10-1996
		AU 5424696 A	02-10-1996
		AU 5424796 A	02-10-1996
		AU 5424896 A	02-10-1996
		AU 5424996 A	02-10-1996
		AU 5425096 A	02-10-1996
		AU 5425296 A	02-10-1996
		WO 9628903 A	19-09-1996
		WO 9628904 A	19-09-1996
		WO 9628905 A	19-09-1996
		WO 9628906 A	19-09-1996
		WO 9628908 A	19-09-1996
		WO 9628909 A	19-09-1996
		WO 9628910 A	19-09-1996
		US 5646942 A	08-07-1997
		US 5559808 A	24-09-1996
		US 5729549 A	17-03-1998
		US 5666365 A	09-09-1997
		US 5852612 A	22-12-1998
		US 5822324 A	13-10-1998
		US 6130898 A	10-10-2000
US 5815794 A	29-09-1998	CA 2184527 A	02-03-1997
		US 5745836 A	28-04-1998
		US 5867764 A	02-02-1999
WO 9748191 A	18-12-1997	AU 3021997 A	07-01-1998
		BR 9709718 A	10-08-1999
		CA 2257967 A	18-12-1997
		CN 1228213 A	08-09-1999
		EP 0904639 A	31-03-1999

Form PCT/ISA/210 (patent family annex) (July 1992)

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No
PCT/US 00/05270

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 5272700 A	21-12-1993	US 5084903 A	28-01-1992
		CA 2079502 A	18-04-1993
		EP 0540188 A	05-05-1993
		JP 6014031 A	21-01-1994
		CA 2009713 A	31-08-1990
		CN 1045318 A,B	12-09-1990
		DE 69025079 D	14-03-1996
		DE 69025079 T	14-11-1996
		EP 0385695 A	05-09-1990
		JP 3034645 A	14-02-1991
US 5825776 A	20-10-1998	AU 711819 B	21-10-1999
		AU 2138497 A	16-09-1997
		BR 9707744 A	27-07-1999
		EP 0883948 A	16-12-1998
		PL 328640 A	15-02-1999
		WO 9732420 A	04-09-1997
US 5373503 A	13-12-1994	NONE	
WO 9839879 A	11-09-1998	US 5991292 A	23-11-1999
		EP 0968586 A	05-01-2000
WO 9815143 A	09-04-1998	FI 963936 A	03-04-1998
		AU 4303497 A	24-04-1998
		EP 0976259 A	02-02-2000
		US 5991639 A	23-11-1999
EP 0790749 A	20-08-1997	NONE	

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
28 December 2000 (28.12.2000)

PCT

(10) International Publication Number
WO 00/79734 A1

(51) International Patent Classification[7]: H04L 12/28, H04Q 11/04

(21) International Application Number: PCT/US00/16913

(22) International Filing Date: 19 June 2000 (19.06.2000)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/139,933 18 June 1999 (18.06.1999) US

(71) Applicant: THE TRUSTEES OF COLUMBIA UNIVERSITY IN THE CITY OF NEW YORK [US/US]; 116th Street and Broadway, New York, NY 10027 (US).

(72) Inventors: DUTTA, Ashutosh; 560 Riverside Drive, Apt. 17C, New York, NY 10027 (US). SCHULZRINNE, Henning; 313 Westview Avenue, Leonia, NJ 07605 (US). YEMINI, Yechiam; 223 Promenade, Edgewater, NJ 07020 (US).

(74) Agents: TANG, Henry et al.; Baker Botts LLP, 30 Rockefeller Plaza, New York, NY 10112-0228 (US).

(81) Designated States *(national)*: AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CR, CU, CZ, DE, DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZW.

(84) Designated States *(regional)*: ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

Published:
— *With international search report.*

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: SYSTEM AND METHOD FOR RECEIVING OVER A NETWORK A BROADCAST FROM A BROADCAST SOURCE

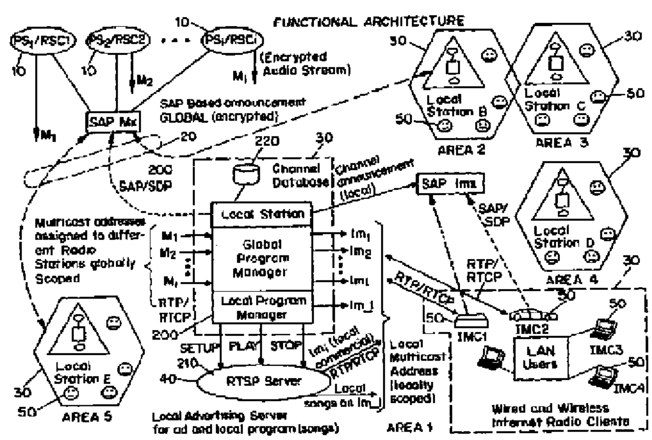

(57) Abstract: A system and method for providing a broadcast to a receiver (IMC 50) via a communication network (20). In particular, the broadcast is received via at least one global multicast channel (Ml). At least one local multicast channel (lm) is associated with the global multicast address (Mx). Then, a communication link is established between the receiver (IMC 50) and the local multicast channel (lm), and the broadcast is routed from the global multicast channel (Ml) to the local multicast channel (lm) to provide the broadcast to the receiver. The number of the receivers which are receiving the broadcast may be determined. The receiver may include an Internet Protocol (IP) interface (IP I/F) which enables the receiver to receive the broadcast via an IP-type multicast communication (20). The receiver may also be wireless (50), and can receive the broadcast in a first subnet (Ia) using a multicast communication. Prior to the receiver (IMC 50) moving to a second subnet (ia), a request is generated by the receiver to receive the broadcast in the second subnet. After receiving the request, the broadcast is provided to the wireless receiver in the second subnet using the multicast communication.

WO 00/79734 A1

SYSTEM AND METHOD FOR RECEIVING OVER A NETWORK A BROADCAST FROM A BROADCAST SOURCE

FIELD OF THE INVENTION:

The present invention relates to a system and method for providing a broadcast over a network to a client. In particular, the system and method utilize network multicast communication for providing the broadcast of content between a broadcast source and the client to avail a global content and/or a local content to user.

APPENDIX

Attached hereto, please find an Appendix which shows an exemplary embodiment of the implementation of the system and method according to the present invention.

BACKGROUND INFORMATION:

Conventional radio systems broadcast a continuous content without requiring extensive user interaction. This traditional scheme is convenient in situations where the listener is sharing his or her attention with other tasks, such as driving an automobile. However, one of the disadvantages of these conventional radio systems is that only a limited number of the radio stations can legally transmit their broadcasts in a particular area (e.g., only 45 FM radio stations can transmit their broadcast in the New York City metropolitan area). There have been a number of proposed solutions to address this limitation. However, none of the proposed solutions effectively utilized the Internet to expand the number of radio broadcasts, as well as television broadcasts, to the wireless users who travel from one geographical area to another.

A streaming real-time multimedia content (which relates to entertainment, music and /or interactive game industries) can now be provided over the Internet. The streaming applications include IP telephony, broadcasting multimedia content and multi-party conferences, collaborations and multi-player

games. However, at least one publication (i.e., the New York Times) asserted that such multimedia streaming applications will bring about the demise of the Internet because the streaming applications are far more demanding in terms of bandwidth, latency and reliability than the traditional data communication applications. Many of the existing streaming systems do not scale to large audiences, particularly for a transmission at high bit rates. They also do not provide a user flexibility, and are restricted to a utilization of either conferencing or broadcast modes.

Early attempts to provide the streaming applications to the clients over the Internet have been implement using a unicast scheme. An exemplary system illustrating the system which utilizes the conventional unicast architecture is shown in Figure 1. Referring to Figure 1, the source 100 (e.g., the audio and/or video content provider) is connected to a first router R1, which in turn is connected to second and third routers R2, R3. The second router R2 is connected to fourth and fifth routers R4, R5, while the third router R3 is connected to sixth and seventh routers R6, R7. The fourth router R4 is connected to two clients C0, C1, the fifth router R5 is connected to three clients C2, C3, C4, the sixth client R6 is connected to two clients C5, C6, and the seventh client R7 is connected to another three clients C7, C8, C9. The clients C0-C9 may be computers requesting the particular multimedia content (e.g., an audio and/or video content).

In operation, if each of the clients C0-C9 requests the same multimedia content, each of those requests is routed via their respective routers to the source 100. Particularly, the clients C0, C1 send such request to the fourth router R4 which routes the request two streams for the particular multimedia content, i.e., one stream for each of its requesting clients C0, C1. At the same time, the fifth, sixth and seventh routers R5, R6, R7 may receive the requests for the same multimedia content from its respective clients C2-C9, and these routers R5, R6, R7 route their streams, respectively, for such multimedia content upstream. The requests for two and three identical multimedia streams (i.e., a total of five streams) are sent to the second router R2 from the fourth and fifth routers R4, R5, respectively. The requests for the same three and two multimedia streams (i.e., also a total of five streams) are sent to the third router R3 from the sixth and seventh routers R6, R7, respectively. The second

and third routers R2, R3 each route the request for five multimedia streams to the first router R1, which routes a request for 10 multimedia streams (i.e., 5 for the second router R2 and 5 for the third router R3) to the source 100.

Thus, the source 100 receives a request for 10 multimedia streams, and then transmits 10 multimedia streams to the first router R1, which then routes the requested 5 identical multimedia streams to the second router R2, and the same 5 multimedia streams to the third router R3. The second router R2 then routes two of these multimedia streams to the fourth router R4, and three to the fifth router R5. The fourth router R4 routes 1 stream to the client C0 and the other stream to the client C1. The fifth router R5 routes one of its received streams to the respective client, C2, C3, C4. Similar routing of the multimedia streams occurs for the third router R3 (and thus for the sixth and seventh routers, (R6, R7).

By utilizing the unicast scheme described above and shown in Figure 1, there may be multiple copies of the same multimedia content being transmitted from the source down to the clients. Such transmission of multiple streams may cause a bottleneck in the network by wasting the Internet bandwidth, and would likely prevent the clients from receiving the multimedia content in an expeditious manner.

Figure 2 shows an arrangement utilizing a conventional multicast communications scheme which addressed at least some of the above-mentioned drawbacks. For the sake of simplicity, the multicast arrangement in Figure 2 is substantially similar to that shown in Figure 1. Using the multicasting communications scheme illustrated in Figure 2, if each of the clients C0-C9 requests the same multimedia content, the routers keep track of the particular client which made the request, and only sends one request for the multimedia stream upstream to the next router in the chain (or to the source 100). For example, the clients C0, C1 may send such request (e.g., a join request) to the fourth router R4, which stores an indication (e.g., a state) therein that at least one of clients C0, C1 sent the particular request. At the same time, the fifth, sixth and seventh routers R5, R6, R7 may receive the requests for the same multimedia content from its respective clients C2-C9, and each these routers R5, R6, R7 stores an indication therein regarding that at least one of their respective clients sent the request for multimedia stream. If the fourth router R4

(or the fifth router R5) already routed the multimedia streams to one of its clients (on the same subnet as the requesting client), it routes the multimedia streams to such requesting client. Otherwise each of the fourth and fifth routers R4, R5 sends a request to receive the multimedia stream that was requested by their respective clients C0-C4 to the second router R2. The second router R2 stores an indication that at least one of the fourth and fifth routers R4, R5 made the request. Each of the sixth and seventh routers R6, R7 also may send a request for the multimedia stream (i.e., that was requested by their respective clients C5-C9) to the third router R3. The third router R3 stores an indication which is similar to the one stored in the second router R2. Then, the second and third routers R2, R3 each send the request for the same multimedia stream to the first router R1, which stores an indication regarding which of the routers R2, R3 made the request. Since the first router R1 is directly connected (or connected in the same subnet) to the source 100, the first router R1 always receives the multimedia stream from the source 100.

In this manner, the first router R1 receives the request, duplicates the received multimedia stream (via multicast channels 500) and transmits 1 copy thereof to each of the second and third routers R2, R3 (if both made the request). The second router R2 then duplicates the received multimedia stream provided in the multicast channels 500, and sends one copy of the stream to each of the fourth and fifth router R4, R5. The fourth router R4, in turn, provides one copy of the received multimedia stream provided in the multicast channels 500 to the client C0 and the other copy to the client C1 (if both made the request). The fifth router R5 duplicates the received multimedia stream, and sends one copy of the received multimedia stream provided by the multicast channels 500 to each of the respective client C2, C3, C4 (if each of theses clients made the request). A similar transmission of the multimedia streams occurs for the third router R3 (and thus for the sixth and seventh routers R6, R7).

With this multicast scheme, the source 100 needs to only transmit one multimedia stream to the requesting router, which in turn duplicates the multimedia stream (if necessary) and transmits a single stream downstream to the routers and/or the clients requesting such stream. Indeed, each router (as well as the source 100)

does not need to transmit more than one multimedia stream to the downstream routers. As such, the bandwidth of the system is utilized more efficiently.

In addition, by using the multicast scheme described above, it is also possible to avoid a transmission of a request for the multimedia stream (that has already been provided to other clients by a particular router) upstream, all the way up to the server 100. For example, another client C10 may be connected to the fourth router R5, and this new client C10 may request the multimedia stream from the fourth router R4 that has already been requested (and is provided to) the client C1. When the fourth router R4 receives this request from the new client C10, it checks whether the requested multimedia stream has already been provided to it. If not, this request is then passed to the second router R2. If the fourth router R4 determines that the requested multimedia stream is already provided by it to at least one of its clients (is in the present exemplary case to the client C1), the fourth router sends a copy of the requested multimedia stream to the new client C10 without sending additional requests for this multimedia stream to the second router R2, and ultimately to the server. Even though this multicast communications scheme provides an advantageous transmission of the multimedia streams from the servers to the clients, it was not effectively usable for wireless communication or in systems where the broadcast streams from different sources which can immediately be provided to the wired or wireless clients.

Previous attempts to provide next-generation radio and television systems have not been successful largely because these systems did not add significant benefits over the older and well known systems. Current versions of the Internet (or web) radio or television were not designed to utilize a large-scale multicast scheme, while also lacking the ability to support low-latency constraints and flexible programming (e.g., an automatic ad insertion during a program, an on-line monitoring of a particular channel, etc.). Furthermore, the conventional systems do not support a continuous streaming or conferencing, while the wireless client is moving, especially from one subnet to another.

SUMMARY OF THE INVENTION

A system and method according to the present invention is provided for transmitting and receiving broadcasts between a broadcast source and a client. One of the exemplary embodiments of the system and method utilizes the available Internet standards and protocols (e.g., RTP, RTCP, RTSP, SIP, SAP, SDP, UDP and IP multicast) to maximize their deployability. Other embodiments of the present invention utilize non-conventional technologies and/or protocols, such as a mobility-aware multicast scheme, a streaming protocol for wireless clients, a fast re-configuration, a bandwidth control for a multicast stream in a wireless network, etc. With the present invention, users can choose to tune-in to receive a local broadcast transmitted by a local station, a global broadcast transmitted by a global station.

The system and method according to the present invention can send broadcasts in a single area, as well as to multiple regions, where there are listeners/viewers who would like to receive the broadcast. This system and method also provides the ability for the end user to invite another user to a particular program using SIP (Session Initiation Protocol). Thus, with the present invention it is now possible to provide:

- Scalable mechanism for a selective content distribution with an automatic localized information insertion by using a hierarchical scope-based multicasting (e.g., global/local multicasting scheme) and local servers.
- Application-layer multicasting arrangement for the real-time broadcast traffic.
- Scalable hierarchical directory structure for an itemized content distribution.
- Support for global and local programs with possible ways of mixing the two.
- Popularity-based spectrum management to address the limits if the spectrum (e.g., a control mechanism for managing an audio/video stream based on a popularity of a particular program - capable of increasing the bandwidth of the broadcast which provides content for broadcasts which are popular with the users).
- Secure payment scheme between the content providers, advertisers and affiliates, which may be utilized for E-commerce.

- Support of a fast-handoff of the Internet Protocol multicast streams when the mobile clients move from one domain to another (e.g., moving in a car on a highway from one subnet to another) in a wireless environment. An application layer mobility protocol and a faster reconfiguration methodology can be provided for the wireless clients to implement such support.
- Distribution of a streaming content to the IP enabled wireless handset (e.g., IP enabled radio/television) using systems with wireless interface and a tuner.
- A combination of intra-ISP multicast with non-multicast global domain (e.g., the unicast domain).
- Support of IP multicast scheme for streaming (e.g., using the MP3 standard) over the bandwidth constrained wireless medium.
- Secure multicast environment to protect against malicious data senders.

One of the embodiments of the system of the present invention provides an architecture to facilitate an IP-based radio/television network, e.g., a streaming network. It can utilize the conventional Internet protocol suite to provide robust communication over conventional heterogeneous access networks. For example, the system and method can also utilize any wired and/or wireless layer-2 technology such as, e.g., PPP ("point to point protocol"), CDMA ("code division multiple access"), protocol based on IEEE 802.11 standard, DSL ("digital subscriber link") and Gigabit Ethernet. It is also possible to utilize the system and method of the present invention other network technologies. The local servers used in the system and method according to the present invention, as well as the use of application layer, provide an degree of scalability. The flexibility of radio services a better reach and a quality of service for the audio/video stream carried over IP are just a few of the other advantageous features of the system and method according to the present invention. Both wired and wireless links may be used for interconnection to the system and method of the present invention, as well as to include various throughput, delay, and error rates. The present invention provides flexible radio/television streaming services to the local Internet (e.g., multimedia clients which may not necessarily be supported by the traditional AM/FM or television receivers). The system and method of the

present invention also provides the flexibility to the clients to be able to receive broadcast from any radio or television station in the world. It offers the capability of a hierarchical searching in terms of categories, and a way to insert local advertisements during commercial breaks. This will meet the challenge of bringing quality audio/video broadcast to the people in remote site, and to the wireless mobile clients. Radio Antenna Servers are provided in the local domains act as local stations/localized servers so as to determine how many people can listen to a particular radio/television station globally without a possible degradation of stream quality and provides the ability for the local listeners in a single domain to switch between the local program and the global program. These servers also provide the ability for the local listeners to receive the local advertisements during commercial breaks, while still being tuned to the global program or to continue listening to a particular segment of the global program while still being tuned to the local program. Another advantageous feature of the present invention is that the system and method allow any server connected to a communications network to be a potential broadcaster. The system and method also provides a pricing model which allows the servers (and possibly the broadcasters) to obtain a direct financial benefit therefrom.

As indicated above, the system according to the present invention is preferably transport independent, operates over wired and wireless links, and accommodates the mobility of the client. Therefore, the present invention provides a continuity to the listener of a particular program broadcast by the local or global station as the mobile client moves. The system and method according to the present invention can also utilize a network topology of highly malleable meshes which would include more than just static trees where each client (or node) can be mobile.

In an exemplary embodiment of the present invention, a broadcast is provided to a receiver via a communication network. The broadcast is received via at least one global multicast channel. At least one local multicast channel is associated with the global multicast address. A communication link is then established between the receiver and the local multicast channel, and the broadcast is routed from the global multicast channel to the local multicast channel to provide the broadcast to the receiver. The number of the receivers which are receiving the broadcast may also be

determined. The receiver may include an Internet Protocol (IP) interface which enables the receiver to receive the broadcast via an IP-type multicast communication. The receiver may also be wireless, and can receive the broadcast in a first subnet using a multicast communication. Prior to the receiver moving to a second subnet, a request is generated by the receiver to receive the broadcast in the second subnet. After receiving the request, the broadcast is provided to the wireless receiver in the second subnet using the multicast communication.

The present invention will will now be described by way of detailed description of exemplary embodiments thereby with reference to the drawings, in which:

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 is a high level functional diagram showing a network based broadcasting system which utilizes a conventional unicast communication scheme;

Figure 2 is a high level function diagram showing a network based broadcasting system of Figure 1 utilizing a conventional multicast communication scheme;

Figure 3 is a functional block diagram showing an exemplary embodiment of a system according to the present invention which utilizes the multicast communication scheme for transmitting and receiving broadcast streams between a source and a client.

Figure 4 is a functional system diagram showing an exemplary implementation of the system illustrated in Figure 3;

Figure 5 is a diagram providing a detailed illustration of the functional architecture of another exemplary implementation of the system of Figure 3;

Figure 6A is a functional block diagram showing an exemplary embodiment of the Internet-capable broadcast receiving devices according to the present invention;

Figure 6B is a functional block diagram showing an exemplary protocol stack, that can be used by the system and method of the present invention;

Figure 7 is a flow diagram representing an exemplary embodiment of the method according to the present invention;

Figure 8 is a flow diagram representing another exemplary embodiment of the method according to the present invention;

Figure 9 is a schematic system-level functional diagram showing a detailed implementation of the system and method according to the present invention utilizing particular protocols;

Figure 10 is a schematic system-level functional diagram showing an exemplary scheme in which multicast systems are interconnected via a non-multicast network;

Figure 11A is a functional diagram illustrating one embodiment of the system and method of the present invention for mobile clients; and

Figure 11B is a functional diagram illustrating another embodiment of the system and method of the present invention for the mobile clients.

DETAILED DESCRIPTION

A. SYSTEM ARCHITECTURE

An exemplary embodiment of the system according to the present invention is shown in Figure 3. The illustrated exemplary embodiment includes four functional components, i.e., a Radio Station Client (RSC) 10 or a Primary Station, a Radio Antenna Server (RAS) 30 or a local station, an Advertisement/Media Arrangement (AMA) 40 and at least one Internet Multimedia Client (IMC) 50. It should be understood that RSC 10 can be a television station client, and RAS 30 can be a television antenna server. IMC 50 can be a car radio or another reception unit which is capable of receiving a multicast broadcast. Such car radio may be an Internet-capable Radio as shall be described in further detail below. In operation, RSC 10 (e.g., a computing device with IP interface) transmits a global multimedia broadcast via a communications network 20 (e.g., the Internet). RAS 30 (e.g., also a server) can receive the global broadcast from the communications network 20, and make this broadcast available to IMC 50 using the multicast communication scheme

described above with reference to Figure 2 and as shall be described in further detail below. In addition, RAS 30 can broadcast a local broadcast to IMC 50, preferably also using the multicast communications scheme as shall be described below. AMA 40 is coupled to RAS 30 so as to insert additional content, indicating advertisements, into the particular segments of the global broadcast that is received from RSC 10 via the communications network 20. AMA 40 can be a separate server with its own storage database or a media database which is within RAS 30. IMC 50 can be used to receive the global broadcast (which may include additional content inserted by AMA 40) as well as a local broadcast by RAS 30.

An exemplary implementation of the system according to the present invention is shown in Figure 4. In this implementation, RSC 10 may include a content server 105. The server 105 (via an Internet Protocol communication arrangement 120) transmits the global broadcast (e.g., the multimedia content) to an arrangement of routers 140 which are part of the Internet (i.e., the communications arrangement 20). These routers 140 deliver the global broadcast to a local station 150 (e.g., part of RAS 30), which can pass this global broadcast to IMC 50. The multimedia content may also be distributed via one or more broadband low earth orbiting satellites 110 to RAS 30, via an earth station arrangement 130. As indicated above, the local station 150 can also provide its own local broadcast to IMC 50. The exemplary implementation shown in Figure 4 preferably utilizes the multicast communication throughout the system. However, if particular portions of the system are not capable of using such multicast communication, it is possible to utilize an alternate scheme in those particular portions as described in greater detail below. It is preferable to implement the multicast communication scheme described above with reference to Figures 2 and 4 between RSC 10 and RAS 30 as well as between RAS 30 and each IMC 50.

Figure 5 shows a detailed illustration of another implementation of the system of Figure 3. This illustration and the illustration provided in Figure 2 shall be referred to below to explain a particular utilization of the multicast communication scheme and how such scheme may be modified in accordance with the system and method of the present invention. In particular, all RSCs 10 have access to a plurality of multicast channels 500 (i.e., addressed at locations M1 to Mi). These addresses 10

may be provided in memory or on the hard drive of one of RSCs 10, in a shared memory distributed between, or may be located on a storage device remote from RSCs 10. The multicast address can also be assigned by a multicast address dispersing computer. In addition, all RSCs 10 have access to a global index address Mx.

In general, a particular one of RSCs 10 may provide a multimedia stream at a particular multicast channel address (e.g., M1), and then announce to the global index address Mx that it has provided the multimedia stream on that particular address. As shall be explained in further detail below, the global multicast addresses are associated with local multicast addresses so that each RAS 40 can forward either the global broadcast provided in at least one of the multicast channels 500 (see Figure 2) broadcast by one or more of RSCs 10, as well as transmit the local broadcast that it generates.

At boot-up time, the clients C0-C9 (i.e, IMCs 50) receive the information associated with the content provided in one or more of the multicast channels 500 (preferably by checking a local index address lmx which is associated with the global index address Mx as shall be described in further detail below). In particular, by checking an address which is associated with the global index address Mx, the clients C0-C9 may determine which multimedia stream is currently being provided in the local channels that are associated, at least in part, with the multicast channels M1-Mi. Then, one or more of RASs 30 may generate the respective requests to receive one or more of the global multimedia streams (provided in the channels which may be associated with the multicast channels M1-Mi). It is also possible for the clients (i.e., IMCs 50) to receive the addresses of the updated multicast channels 500 from the source (i.e., RSC 10) in real-time or when desired. The requests are transmitted upstream to the routers (not shown) which are connected to the respective clients (i.e., IMCs 50).

Provided below is a detailed description of the exemplary components of the illustrative system and method according to the present invention described above, with reference to Figure 5.

I. Radio/Television Station Client (RSC)/Primary Station

As indicated above, RSC 10 can be a computing device of any regular radio/television station/broadcaster that is capable of transmitting its regular programming on an Internet Protocol-based network. It should be understood that Radio Station Client (RSC) can also be a station client which transmits a television type broadcast over the communications network. When RSC 10 broadcasts its program over the communications network 20 (e.g., the Internet), such broadcast is transmitted to an Internet gateway (not shown in Figure 5) (e.g., a router) located near the server's location. Each primary station of RSC 10 (e.g., PS1, PS2 ... PSn as shown in Figure 5) can preferably transmit its broadcast on an assigned unique multicast channel corresponding to a particular multicast address (e.g., M1, M2... Mi), and the respective broadcasted content is provided to this address. As discussed above, the assigned multicast address, along with few other relevant parameters, are announced to a global multicast address (Mx).

II. Antenna Server(RAS)/Local Station

RASs 30 are generally distributed according to the population, the geographic area and/or some other topology. Each RAS 30 preferably offers two program tracks to a user of IMC 50 - the global broadcast transmitted by RSC 10 and the local broadcast provided by RSC 30. In should be understood that RAS 30 can transmit/receive televison broadcasts. Since numerous global broadcast can be provided on a number of multicast channels, RSC 30 preferably relays at least a subset of all transmitted programs in the global broadcast to IMC 50. The broadcast transmitted by RSC 10 is generally transmitted globally with gaps in the global broadcast so that the local advertisement and/or promotional content can be inserted in such gaps. The local broadcast may be local news segments provided by RAS 30. This scheme according to the present invention provides the user of IMC 50 with an ability to receive either the local broadcast or the global broadcast.

RAS 30 preferably includes a Management Server (MS) 200 and a channel database 220. The Management Server 200 creates and/or maintains the channel database 220, records the statistics regarding the number of IMCs 50 that are

receiving a particular broadcast at a particular local multicast channel, provides control tools for maintaining and modifying configurable parameters, and manages the interface with other devices (e.g., a RTSP server and/or media database, etc.). For each RAS 30, the Management Server 200 monitors the global index address Mx, and receives the global multicast channels M1, M2 ... Mi (which provide the audio and/or video streams) that are described by the global index address Mx.

These multicast channels are provided in an encrypted form to RAS 30. An exemplary scheme to decrypt the encrypted multicast channels at RAS 30 shall be described in further detail below. After decrypting one or more of the global multicast channels M1, M2 ... Mi, the stream provided at the address of the decrypted multicast channel (e.g., the global channel M1) is rerouted to a particular local multicast channel (e.g., the local channel lm2) that is provided at a corresponding local address. In this manner, IMCs 50 can receive the decrypted stream which is provided at the global channel M1 to RAS 30. RAS 30 also maintains the directory services, and keeps track of the IMCs 50 that receive a particular broadcast (i.e., local and/or global). Hence, RAS 30 can provide pay-per-listen and/or pay-per-view channels, bill the subscriber using the IMCs 50 and manage them.

III. Advertisement/Media Arrangement (AMA)

As described above, RAS 30 may include AMA 40, or AMA 40 can be provided remotely from RAS 30. AMA 40 includes a Local Advertisement Server 210 (which can be an RTSP server). This Local Advertisement Server 210 is capable of playing local media on demand programs (e.g., songs and/or music videos), as well as inserting a local advertisement into the global broadcast during a commercial break thereof.

IV. Internet Multimedia Client (IMC):

IMC 50 can be a wired Internet Protocol (IP) device or a wireless IP device. For example, IMC 50 can be considered wired when it is connected on a LAN, and wireless when it is located remote from the LAN and communicating over a wireless communications link. IMC 50 is capable of executing application programs

which monitor the local index multicast address lmx where data regarding the global or local program are provided. Conventional tools (e.g., NeVot, Vic, vat or any tool based on SAP/SDP standards) can be utilized by IMCs 50 to monitor the broadcasts and receive the multimedia (e.g., audio and video) streams from the local multicast channels lm1, lm2 ... lmi. Using these tools, IMCs 50 may select any of the broadcasts (i.e., local or global) provided by RAS 30 by e.g., viewing the local multicast index address lmx on the displays of IMCs 50.

Once, IMC 50 selects a particular channel, it starts sending an RTCP signal and receives the audio and/or video stream over UDP/IP. The protocols described herein (e.g., RTPC, UDP/IP, etc.) are known in the art, some of which shall be described below in a greater detail. After receiving the RTCP signal, the Management Server 200 starts monitoring the global multicast address of the global multicast channel which provides the broadcast (e.g., the radio program) selected by IMC 50. When the broadcast at the selected channel is detected, RAS 30 directs it to the assigned address of the local multicast channels. The Management Server 200 continues to transmit the broadcast content, and only interrupts the broadcast when there are no more IMCs 50 that are receiving and/or requesting this broadcast.

As shown in Figure 6A, IMC 50 can be a radio having an ability to toggle between AM/FM broadcasts and the Internet channels, and/ or a television which can receive wireless and/or cable broadcasts, as well IP broadcasts. For example, it is possible to provide a wireless interface having UDP/IP multicast stack which can be connected to a conventional portable radio or a portable television, (or utilized independently). Thus, the connection of an conventional radio/television receiver to the Internet can be accomplished. As an example, the conventional radio/television receiver includes a tuner for AM/FM broadcasts and/or for the television broadcasts. In addition, this radio/television receiver may include a switch (e.g., a mechanical switch, an electrical switch, an automatic software switch, etc.) with which the radio/television receiver can be converted to an Internet-ready device. Based on the SDP parameters of the program being broadcasted, the tuner of the Internet-ready device would detect the broadcasts and possibly categorized them (e.g., News, Entertainment, etc.). Advantageously, the categories and the available

broadcasts are presented on a display screen of such device so that the user can select which category/broadcast he or she would like to receive.

It is also possible to utilize a conventional speech generation/recognition system in connection with the Internet-ready device. For example, the device would provide the available broadcasts/categories to the speech generation/recognition system which would then generate voice-type descriptions of the broadcasts/categories. Then, the user may vocalize his or her selection, and the speech generation/recognition system would determine the selection and provide the requested action.

B. EXEMPLARY PROTOCOLS AND OPERATION/IMPLEMENTATION

I. Protocols

The system and method according to the present invention uses (and possibly modifies) the conventional protocols, i.e., SAP (Session Announcement Protocol), SDP (Session Description Protocol), RTSP (Real-Time Streaming Protocol), RTP (Real-time Transport Protocol), TCP, UDP, IP and IP Multicast. An exemplary protocol stack utilized by the exemplary embodiment of the system and method is shown in Figure 6B. The network infrastructure can be wired and/or wireless. One exemplary implementation of this infrastructure can operate with LMS/MMD wireless links.

Provided below is a short description of the primary protocols that can be used by the exemplary embodiment of the system and method of the present invention.

SDP is a Session Description Protocol which is usable for multi-media sessions, and can be utilized as a format for a session description (generally does not incorporate a transport protocol). SDP is intended to be used for different transport protocols as appropriate, including SAP, SIP, RTSP, electronic mail using MIME extensions, and HTTP. SDP includes the session name and purpose, the time the session is alive, the content type (e.g., audio and/or video) comprising the session, information to enable reception of those content types (addresses, ports, formats etc.), the bandwidth to be used by the broadcast, and the contact information for the person

responsible for session. SDP is widely used for the multicast sessions over the Internet. In order to assist in the advertisement of multicast sessions and to communicate relevant session setup information to prospective participants, a distributed session directory can be used. An instance of such a session directory periodically multicast packets containing a description of a multimedia session to a multicast address. These signals are subsequently received by potential participants, who can use the session description to start the tools required to participate in the session. Using this protocol, the sender can assign a particular bandwidth for a particular application (e.g., radio and/or television broadcast). In this manner, the more popular or bandwidth-intensive application (e.g., television news) would use more bandwidth than non-popular application/broadcast. Thus, a popularity-based spectrum management can be achieved.

SAP is an announcement protocol that distributes the session directory to the multicast conference sessions. An SDP datagram is part of the payload for SAP. SAP client which announces a conference session, periodically multicasts an announcement packet to a known multicast address and port. The appropriate address is determined by the scope mechanisms operating at the sites of the intended participants. IP multicast sessions can be either TTL-scoped or administratively scoped. Thus, an instance of the session directory may need to listen on multiple multicast addresses. The announcement contains a session description and optionally an authentication header. The session description may be encrypted. It is preferable to provide an authentication and integrity of the session announcements to ensure that only authorized parties modify session announcements, and to provide the facilities for announcing the securely encrypted sessions while providing the relevant proposed conferees with the means to decrypt the data streams.

RTSP is a client-server multimedia presentation control protocol which is used for an efficient delivery of streamed multimedia over IP networks. It utilizes the existing web infrastructure (e.g., inheriting authentication and PICS from HTTP). This application level protocol may provide the robust streaming multimedia in one-to-many applications via unicast and multicast communication arrangements, and may support the interoperability between the clients and the servers from different

vendors. The process of streaming breaks media streams into many packets sized appropriately for the bandwidth available between the client and the server. When the client receives enough packets, the user software can be playing one packet, decompressing another, and receiving a third. The user can begin listening almost immediately without the necessity to download the entire media file. RTSP can control multiple data delivery sessions, and is capable of providing a way for selecting the delivery channels (such as UDP, TCP, IP Multicast) and delivery mechanisms based on RTP. RTSP can be used in conjunction with other protocols to set up and manage the reserved-bandwidth streaming sessions.

RTP is a thin protocol which provides support for applications with real-time properties which can be run over UDP. RTP provides a timing reconstruction, loss detection, security and content identification. RTP can be used, possibly without RTCP, in the unicast or multicast communication arrangements. In order to set up an RTP session, the application may define a particular pair of the destination transport addresses (e.g., one network address and a pair of ports for RTP and RTCP). In a multimedia session, each medium (e.g., audio, video, etc.) can be transported in a separate RTP session with a corresponding RTCP session reporting the reception quality.

RTCP may operate in conjunction with RTP. It provides support for the real-time conferencing of large groups on the Internet. RTCP control packets are periodically transmitted by each participant in an RTP session to all other participants. The feedback of the information to the application can be used to control the performance and for other diagnostic purposes. RTCP provides the following exemplary functions:

Feedback to sending application regarding the quality of the data distribution.

Identification of the RTP source.

RTCP transmission interval control.

Communication of the minimal session control information.

SIP has been adopted by the industry, in many cases, as the signaling protocol for the Internet conferencing and telephony. SIP is a client-server protocol

which provides the mechanisms so that the end systems and the proxy servers can provide different required services for setting up a proper signaling scheme. SIP creates, modifies and terminates the associations between the Internet systems (e.g., conferences and point-to-point calls). SIP is a text-based protocol similar to HTTP and RTSP, in which the requests are issued by the client, and the responses are returned by the server. SIP is independent of the packet layer and only utilizes a datagram service, since it provides its own reliability mechanism. This "light-weight" protocol is typically used over UDP or TCP, and provides light-weight signaling. SIP supports the unicast and multicast communication schemes, as well as combinations of thereof. It can implement a variety of the conference-related services with a small set of handling primitives.

II. EXEMPLARY IMPLEMENTATION USING THE PROTOCOLS

The general implementation of an exemplary embodiment the system and method according to the present invention has been already described above. An exemplary implementation of the system and method utilizing the above-discussed protocols is as follows.

a. Channel Announcement

With reference to Figure 5, according to the present invention, a particular RSC 10 may send its program live on a unique global multicast channel (e.g., M1) globally scoped and encrypted using RTP/UDP. Other RSCs 10 can also broadcast their programs on other global multicast channels. Indeed, the multicast channel address is different for each broadcast and/or for each RSC 10. These stations send their session announcement using a subset of SDP parameters to the global index multicast address Mx (which can be encrypted). This common global multicast address contains a list of the programs that are being broadcasted by RSCs 10 on the communication network 20. SDP or a variant thereof can be modified to provide IMCs 50 with additional details regarding the streaming being broadcasted.

b. Channel Management

Figure 7 shows a flow diagram representing an exemplary implementation of one embodiment of the method according to the present invention. In particular, each RAS 30 has a global encryption key which is used by the respective RAS 30 to monitor the global index multicast address (Mx) to obtain, e.g., the listing of the channels and the contents of the channels (step 300). Then, it is determined (e.g., using a decryption technique) if RAS 30 can receive some or all global broadcasts (step 310). If so, RAS 30 is then provided with an authorization to utilize the global broadcast on the global multicast channels M1 ...Mi provided by RSC 10 (step 320). Either automatically or via the manual control, RAS 30 may decide to broadcast at least a part of the list to IMCs 50 that are associated with RAS 30. For this purpose, RAS 30 may create and/or utilize the channel database 220 which contains the list of the supported channels, each with their appropriate attributes, to associate the global broadcast channels with the local broadcast channels (step 330). The subset of channel descriptions announced by each RSC 10 provides sufficient data for generating and updating this database 220, which may be a subset of the list that is received from the global index multicast address Mx. In this manner, the association between the global and local multicast channels can be recorded in the channel database 220 (step 340).

Then, it is determined if RAS 30 is also transmitting a local broadcast (step 350). If so, RAS 30 transmits its local programs on a specific local multicast address lm_l, and records this information in the channel database 220 (step 360). If it is determined in step 350 that RAS 30 is not transmitting the local broadcast, the process proceeds to step 370, in which RAS 30 either generates and/or modifies the information in the channel database 220 regarding the broadcasts (e.g., local and/or global broadcasts) which are available for IMC 50. In step 380, RAS 30 sends the information provided on the local index multicast address lmx for the announcement using SAP to its IMCs 50. RAS 30 also sends the announcement regarding its own local programs to the same local index multicast address lmx using SAP. The announcement on the local index multicast address lmx is preferably not encrypted since the RAS 30 prefers all its clients (i.e., the associated IMCs 50) to see what is

being broadcasted by it. In an alternative exemplary embodiment of the method of the present invention, RAS 30 maintains a pair of multicast addresses for each channel to maintain an association between the global multicast channel address (e.g., M1). Using the respective channels, RSC 10 provides its global program on the local multicast channel address (e.g., lm2) on which the broadcast being is transmitted to IMCs 50 by RAS 30 (i.e., steps 330 and 340).

Figure 8 shows a flow diagram representing yet another exemplary embodiment of the method according to the present invention which is executed when the information in the local index address lmx is provided to IMC 50. In particular, IMC 50 receives the information in this local index address lmx (step 400). Then, in step 410, IMC 50 may request to receive the broadcast from a particular local multicast channel (e.g., lm2). This broadcast can be encrypted or un-encrypted depending on the type of a payment model being utilized. Then, RAS 30 determines, based on the information regarding the local multicast address being requested by IMC 50, whether the broadcast on the particular channel is local or global (step 420). If it is determined that the requested broadcasted is a global broadcast (i.e., originated from RSC 10), RAS 30 uses the channel database 220 to route the global broadcast from the global multicast channel on which the requested broadcast is being transmitted to a corresponding local multicast address (step 430), and the process is directed to step 450. IMC 50 continues transmitting the RTCP packets to the Management Server 200 of RAS 30 as long as it receives the global broadcast on the particular local multicast channel. It should also be noted that when the Management Server 220 receives the global broadcast from RSC 10 on a specified multicast address using RTP/UDP, it also periodically exchanges RTCP signals with RSC 10.

If it is determined that the requested broadcast is a local broadcast (i.e., originated from RAS 30), RAS 30 provides the local broadcast to IMC 50 on the local multicast channel lm_l which is assigned for local broadcasts (step 440). If RAS 30 indicates that another broadcast (either pre-recorded or live) or an advertisement should be inserted into the global or local broadcast (step 450), RAS 30 inserts (or plays) such broadcast and/or advertisement into the local multicast channel associated with the local multicast address of the global or local broadcasts address using, e.g.,

SETUP and PLAY commands (step 460). For example, the inserted broadcast may be either a live news broadcast or a prerecorded news broadcast. Then, RAS 30 provides the requested broadcast on the corresponding local multicast broadcast channel (e.g., lm2), either with or without the additional content being inserted into the broadcast (step 470). Thus, for that particular period, a local manager of RAS 30 may decide to join such specific global multicast group, this may be done when the local manager receives the RTP packets from RSC 10, and generates the RTP/RTCP packets for IMC 50 on the respective local multicast address.

c. Using the Protocols

To summarize, IMCs 50 may be Internet Multimedia Clients (e.g., personal computers and laptops utilizing wired and/or wireless interconnect, car radios/televisions having the IP interface) which monitor the local index multicast address lmx to determine what is available. Such monitoring can be performed using SAP- and/or SDP-based tools. As described in the SAP specification (which is incorporated herein by reference) and as known to those having ordinary skill in the art, RAS 30 can update the announcement information approximately every few minutes. Thus, the program executed at IMC 50 may wait for few minutes before seeing the most updated channel information. By using SAP, this lag is either substantially reduced, or even eliminated, by a caching scheme. For example, this caching scheme either executes the SAP receiver of IMC 50 in the background to continuously keep its cache current, or moves to a local SAP proxy at the startup time of IMC 50 and requests a cache download. In the latter case, RAS 30 essentially becomes the SAP proxy.

When IMC 50 makes a request to listen to one of the programs listed in the program listing (e.g., clicks on the channel), this IMC 50 sends the RTCP signal to the local station manager of RAS 30. If there is a broadcast (e.g., a data stream) already playing on this local multicast address provided pursuant to a previous request from other IMCs 50, then this particular IMC 50 starts receiving the audio and/or video stream using RTP/UDP. However, if this is the first request for such broadcast in this local domain, then RAS joins the multicast tree of the corresponding global

multicast address to receive the broadcast from the corresponding RSC 10 which is transmitting the requested broadcast. RAS 30 can use a conventional application program (e.g., "mlisten") to determine if there is any member which is part of any particular multicast group that is currently transmitting broadcasts, and thus should be able to determine if the request is a first such request for a particular multicast group. "mlisten" is a conventional multicast application for monitoring the number of users joining a particular multicast group (e.g. receiving information from a particular multicast channel).

d. Local Advertisement Insertion

In accordance with exemplary embodiments of the system and method of the present invention, the insertion of advertisement content into the global or local broadcast transmitted to IMC 50 is now described below. The system is implemented such that RSC 10 knows the starting time and the duration of a commercial break prior to the transmission of the global broadcast, since it controls the time for such break. These commercial breaks can also be event driven. Along with the RTP packets, RSC 10 continues sending the RTCP packets to the global multicast address of the global multicast channel where RASs 30 are monitoring the streams of broadcasts. Using the RTCP report, RSC 10 provides the signal to RAS 30 which indicates the time and the duration of a break in the broadcast. The term "advertisement" as used herein includes not only the content directed to selling a product or service or to promote the goodwill of a commercial sponsor, but also to public service messages and announcements, station break announcements, promotions and/or other programming to be broadcasted.

Upon receiving such signal, the Management Server 200 of RAS 30 requests the local RTSP server 210 (which is part of AMA 40) to start playing the local advertisement from a storage medium to a specific local multicast address which is associated with the global multicast address at which RSC 10 transmits the global broadcast. RAS 30 uses a set of RTSP commands, such as SETUP, PLAY and STOP on AMA 40. During this time, the Management Server may stop forwarding the RTP stream from the global multicast channel to the associated local multicast channel.

The local advertisement runs for a time determined by the Management Server 200 using the information received from the RTCP reports. At the end of the time for the commercial break, the Management Server 200 sends a STOP signal to the RTSP server 210 so that it stops playing on that particular multicast address. Then, the Management Server 200 resumes redirecting the audio and/or video streams from the global multicast address to the associated local scoped multicast address. Since the commercial break times for RASs 30 may overlap, it is possible that the RTSP server 210 could play several different local advertisements on the different local multicast addresses. An illustration of the exemplary implementation described above is shown in Figure 9.

One implementation of the system and method for inserting the advertisements into the broadcasts is described in greater detail below. In particular, the system and method can use "InsertAd.java" commands to insert local advertisements. As soon as the broadcast appears on the global multicast channel, it starts an InsertAd thread which listens for RTCP packets generated by RSC 10 (e.g., the RTCP port is one greater than the RTP port).The RTCP packets from RSC indicate the number of seconds remaining until the start of the advertisement, as well as the length thereof. InsertAd command inserts a local advertisement by switching the channel mode to "advertisement". When the global commercial is finished, InsertAd switches the channel mode back to "redirect". The list of the local advertisement files is specified inside a list file which are inserted using, e.g., a Round Robin scheduling scheme.

At startup, RSC 10 initiates RsSendRTCP thread to notify RAS 30 of the commercial breaks. The thread sends the RTCP packets so that RAS 30 can insert local advertisement. The RTCP packets indicate the number of seconds remaining to the start of the advertisement, as well as the length of the advertisement. The start times of the advertisement are read in the following format (e.g., one record per line):

day/hour/minute/second/duration;

day is 1 through 7 which stands for Sunday through Saturday; hour is 0 through 23; minute is 0 through 59; second is 0 through 59; and duration is specified in seconds.

Since the RTCP packets are transmitted over UDP, there may be a possibility of a packet loss or an incorrect order. To address this potential problem, the RTCP packets are re-transmitted (e.g., one packet 4 seconds prior to the advertisement, next one -3 seconds, next one -2 seconds, etc. with the corresponding value in the field which indicates the time remaining until the start of the commercial). In addition, to allow RAS 30 to distinguish between the re-transmissions and advertisements, the RTCP packets have a particular sequence number. All re-transmissions have the same sequence number. Each advertisement has a sequence number one greater than the previous sequence number.

e. User Interface for Itemized Content

It is possible to utilize and/or modify a conventional directory structure/user interface referred to as "sdr" in implementing the embodiments of the system and method of the present invention. This directory structure/user interface can be used as a tuning mechanism for the wireless IP radios and/or IP television, and may be touch-tone based, voice activated, etc. This "sdr" structure/program (created by ISI, Meriana del Ray, CA) can be modified or extended to make it more customized and searchable for searching purposes. For example, "sdr" can be modified to categorize the content of the streaming media according to the type of program being broadcasted (e.g., "game show", "news", etc.) In addition, it is possible to utilize a voice activated-type "sdr" according to the content type, as well as to provide a menu for a particular locality. Also, with a touch of a button or by pronouncing a particular word (e.g., "News"), sdr would provide a visual menu or a voice menu to indicate which channels are available to that particular locality. With another touch tone or voice activation, sdr may provide IMC 50 with access to the broadcast from the local multicast channel. Other features need not be further discussed, since they would be clearly understood to one having ordinary skill in the art.

f. Payment Model

There are numerous payment models that can be supported by the system and method according to the present invention. For example, RAS 30 may

collect the fees from the local advertisement sponsors for broadcasting their advertisements during the commercial breaks while relaying the global or local station broadcasts. In addition, RAS 30 may also relay some pay-per-listen and/or pay-per-view programs. In this case, RAS 30 pays the global station (i.e., RSC 10) a fee which depends on how many listeners/viewers are listening to or viewing a particular program. The number of listeners/viewers can be determined from the RTCP reports that are generated from IMCs 50. Every RAS 30 can also broadcast its local program to IMCs 50 with the segments of the news or some other premium programs relayed from RSCs 10.

A different type of the pricing model can also be provided to reflect the process of determining when and on which channels the advertisers should place their advertisements in order to maximize their return on investment. Priorities can be assigned to certain advertisements so as to enable the advertisers to compete for a higher time slot or timing of the advertisement (e.g., the highest paying company would get the slot during the Super Bowl by using a contention algorithm). It may also be possible to implement the exemplary payment schemes, e.g., public financing, advertising and on-air solicitations for donations. Hybrid models (e.g., the paying customers are not required to view or listen to commercial or receive solicitations for donations) are also feasible. Furthermore, another embodiment of the payment model can be associated with the security model described below.

g. Security

It is possible to provide at least four levels of encryption for the system and method according to the present invention (e.g., a global announcement encryption, a global multicast stream encryption, a local audit encryption and a user authentication).

Utilizing the global announcement encryption, it is possible to separate the global announcements from the local announcements. IMCs 50 should not be able to gain access to the global announcements, and would only be able to view the local announcements. With the global encryption key during the announcement (by RSCs 10), IMCs 50 would not be allowed to find out about available the global channels,